arXiv:1403.3325v1 [math.PR] 13 Mar 2014

Slow transitions, slow mixing, and starvation in dense random-access networks

A. Zocca* S.C. Borst*† J.S.H. van Leeuwaarden*

Abstract

We consider dense wireless random-access networks, modeled as systems of particles with hard-core interaction. The particles represent the network users that try to become active after an exponential back-off time, and stay active for an exponential transmission time. Due to wireless interference, active users prevent other nearby users from simultaneous activity, which we describe as hard-core interaction on a conflict graph. We show that dense networks with aggressive back-off schemes lead to extremely slow transitions between dominant states, and inevitably cause long mixing times and starvation effects.

Keywords: wireless random-access networks; hitting times; throughput analysis; starvation phenomena; mixing times.

1 Introduction

We consider a stylized model for a network of N users sharing a wireless medium according to a random-access scheme. The network is represented by an undirected graph $G = (V, E)$, called *conflict graph*. The set of vertices $V = \{1, \ldots, N\}$ describes the network users and the set of edges $E \subseteq V \times V$ indicates which pairs of users interfere and are thus prevented from simultaneous activity. The independent sets of G (sets of vertices not sharing any edge) then correspond to the feasible joint activity states of the network.

A user is said to be *blocked* whenever the user itself or any of its neighbors in G is active, and *unblocked* otherwise. User i activates (starts a transmission) at an exponential rate ν_i whenever it is unblocked, and then remains active for an exponentially distributed time period with unit mean, before turning inactive again. The durations of the various activity periods are assumed independent across time and among users. We will refer to the parameters ν_i as *activation rates*.

Let $\Omega^* \subseteq \{0,1\}^V$ be the collection of incidence vectors of all independent sets of G, and let $X^*(t) \in \Omega^*$ be the joint activity state at time t, with element i of $X^*(t)$ indicating whether user i is active ($X_i^*(t) = 1$) or not ($X_i^*(t) = 0$) at time t. Then $\{X^*(t)\}_{t\geq 0}$ is a reversible Markov process with stationary distribution [4, 15, 16, 17, 26]

$$\pi_x(\nu_1, \ldots, \nu_N) = \lim_{t\to\infty} \mathbb{P}\left(X^*(t) = x\right) = \frac{\prod_{i\in V} \nu_i^{x_i}}{\sum_{y\in\Omega^*} \prod_{i\in V} \nu_i^{y_i}}, \quad x \in \Omega^*. \tag{1}$$

We also mention that the model amounts to a special instance of a loss network [24, 27], and that the product-form distribution (1) corresponds to the Gibbs measure of the hard-core model in statistical physics [5, 7].

For the case $\nu_i = \nu$ it follows from (1) that only the activity states corresponding to *maximum* independent sets retain probability mass as the activation rate ν grows large. This indicates that users that do not belong to a maximum independent set have far fewer opportunities to be active. This disadvantage is commonly referred as *spatial unfairness*, and the associated starvation effects have major performance repercussions in wireless networks.

It has been shown that spatial unfairness can be avoided by selecting suitable user-specific activation rates ν_i which provide all users with an equal opportunity to be active in the long run [8, 25]. Even in those cases, however, or in symmetric scenarios where spatial fairness is automatically ensured, transient but yet significant starvation effects can arise due to extremely slow transitions between high-likelihood or *dominant* states. Intuitively speaking, the activity process will typically need to pass through a low-likelihood or bottleneck state in order for the process to transit between dominant states. Visiting such a bottleneck state basically involves the occurrence of a rare event, or even several rare events in different limiting regimes, and causes the transition to take a correspondingly long

*Department of Mathematics and Computer Science, Eindhoven University of Technology
†Alcatel-Lucent Bell Labs, Murray Hill, NJ, USA

amount of time. Consequently, users may experience extended stretches of forced inactivity (possibly interspersed by long intervals with a rapid succession of activity periods), resulting in serious performance degradation.

Motivated by these fairness issues, we investigate in the present paper the time for the Markov process to reach, starting from a given dominant state, one of the other dominant states. We study these hitting times as well as mixing properties in the asymptotic regime where the activation rates ν_i grow large. This asymptotic regime, in which users activate aggressively, is relevant in highly loaded networks and gives rise to the above-described starvation effects. As shown numerically in [18], these starvation effects are particularly pronounced in dense topologies. As a prototypical worst-case scenario, we focus on a specific class of dense conflict graphs, namely complete partite graphs. In such networks the users can be partitioned into K disjoint sets called *components*, such that each user interferes with all users in all other components. This implies that a transition from an activity state in one of the components to another component entails passing through a bottleneck state where all users are inactive at some point. Based on this observation and a regenerative argument, we establish a geometric-sum representation for the hitting time, which we then use to obtain the asymptotic order-of-magnitude and scaled distribution. For convenience we assume all users within a given component to have the same activation rate, but we do allow for users in different components to have different activation rates. Section 2 presents a detailed model description and Section 3 gives an overview of the main results. Preliminary results for the case $\nu_i = \nu$ for all $i \in V$ appeared in [28], but did not exhibit the full qualitative range of asymptotic behaviors that will be revealed in the present paper.

2 Model description

Consider a network represented by a *complete partite graph*, where two users interfere if and only if they belong to different components. Thus, in particular, users within the same component do not interfere. Denote by $C_1, \ldots, C_K$ the K components of G and define $L_k := |C_k|$ as the size of component C_k. Note that the components $C_1, \ldots, C_K$ are the K *maximal* independent sets of the graph G. Moreover, component C_k corresponds to a *maximum* independent set if and only if $L_k \geq L_j$ for all $j = 1, \ldots, K$. Figure 1 shows an example of such a dense conflict graph, where $K = 5$ and the components have sizes $\{L_1, L_2, L_3, L_4, L_5\} = \{3, 4, 6, 2, 5\}$. The corresponding state space Ω^* for this graph is shown is Figure 2(a).

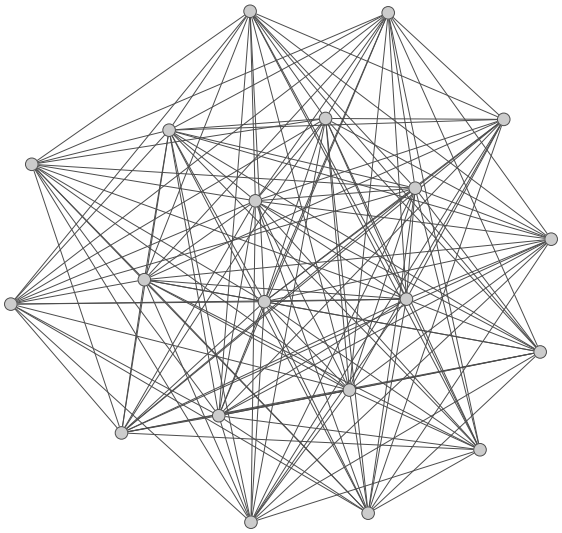


Figure 1: Example of complete K-partite conflict graph with $K = 5$.

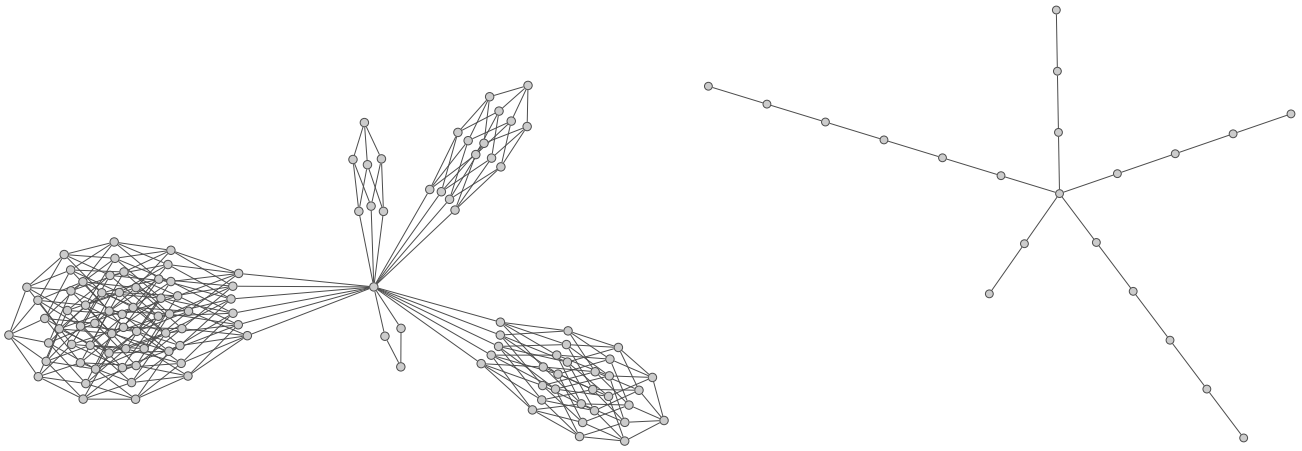
We assume that the exponential rate at which a user activates depends only on a global aggressiveness parameter ν and on the component it belongs to, namely

$$\nu_i = f_k(\nu) \quad \text{if } i \in C_k,$$

for some monotone function $f_k : \mathbb{R}_+ \to \mathbb{R}_+$ with $\lim_{\nu\to\infty} f_k(\nu) = \infty$. We will refer to the function $f_k(\cdot)$ as the *activation rate* of component C_k, for $k = 1, \ldots, K$.

In view of symmetry, all states with the same number of active users in a given component can be aggregated, and we only need to keep track of the number l of active users, if any, and the index k of the component C_k they belong

to. This state aggregation yields a new Markov process $\{X(t)\}_{t\geq 0}$ on a star-shaped state space Ω with K branches, where each branch emanates from a common root node and describes one of the components of the conflict graph. Figure 2(b) shows the aggregated state space corresponding to the previous example.



(a) State space Ω^* (b) Aggregated state space Ω

Figure 2: State space Ω^* and aggregated state space Ω, for the conflict graph in Figure 1.

For $k = 1, \ldots, K$, let

$$\mathcal{B}_k := \{(k, l) : 1 \leq l \leq L_k\}$$

denote the branch of the state space Ω that corresponds to activity inside component C_k, where state (k, l) indicates l users active in component C_k. Then $\Omega = \{0\} \cup \bigcup_{k=1}^{K} \mathcal{B}_k$, where 0 is the bottleneck state in which all users are inactive.

The transition rates of the process $\{X(t)\}_{t\geq 0}$ then read

$$\begin{aligned}
q(0, (k, 1)) &= L_k f_k(\nu), \\
q((k, 1), 0) &= 1, \\
q((k, l), (k, l + 1)) &= (L_k - l) f_k(\nu), \quad l = 1, \ldots, L_k - 1,\ k = 1, \ldots, K, \\
q((k, l), (k, l - 1)) &= l, \quad l = 2, \ldots, L_k,\ k = 1, \ldots, K.
\end{aligned}$$

The stationary distribution of the process $\{X(t)\}_{t\geq 0}$ can be easily seen to be

$$\begin{aligned}
\pi_0(\nu) &= \left(1 + \sum_{k=1}^{K} \sum_{l=1}^{L_k} \binom{L_k}{l} f_k(\nu)^l\right)^{-1}, \\
\pi_{(k,l)}(\nu) &= \pi_0(\nu) \binom{L_k}{l} f_k(\nu)^l, \quad l = 1, \ldots, L_k,\ k = 1, \ldots, K.
\end{aligned} \tag{2}$$

The state (k, L_k) corresponds to the maximum activity state inside component C_k, which becomes the most likely state within the branch $\mathcal{B}_k$ as $\nu \to \infty$. Define the *transition time* from state (k_1, l_1) to state (k_2, l_2) as

$$T_{(k_1,l_1),(k_2,l_2)}(\nu) := \inf\{t > 0 : X(t) = (k_2, l_2) \mid X(0) = (k_1, l_1)\}.$$

We now introduce few parameters that will turn out to play a key role in the asymptotic distribution of the transition time. Define for $k \neq k_2$,

$$\gamma_k := \lim_{\nu\to\infty} \frac{f_k(\nu)^{L_k}}{\sum_{j\neq k_2} f_j(\nu)^{L_j}}. \tag{3}$$

To avoid technicalities, we assume throughout that all parameters γ_k are well defined. In view of (2), γ_k may be interpreted as the stationary fraction of time that the activity process spends in branch k as $\nu \to \infty$, excluding the target branch k_2. As it turns out, γ_k also equals the fraction of time that the activity process spends in branch $\mathcal{B}_k$ during the transition time as $\nu \to \infty$.

Branch $\mathcal{B}_k$ is called *dominant* if $\gamma_k > 0$ and let $K_* := \{k \neq k_2 : \gamma_k > 0\}$ be the index set of all *dominant branches*. Note that, by construction, the set K_* is never empty and thus there is always at least one dominant branch.

3 Main results

In this section we present our main results, which are all related to the asymptotic behavior of the transition time $T_{(k_1,l_1),(k_2,l_2)}(\nu)$ in the asymptotic regime of a large activation rate ν.

Our first result characterizes the asymptotic order-of-magnitude of the mean transition time in terms of the activation rates and the network structure. For any two real-valued functions $f(\cdot)$ and $g(\cdot)$, let $f(\nu) \sim g(\nu)$ indicate that $\lim_{\nu\to\infty} f(\nu)/g(\nu) = 1$ as $\nu \to \infty$.

Theorem 3.1. *If $k_1 \neq k_2$, then*

$$\mathbb{E}T_{(k_1,l_1),(k_2,l_2)}(\nu) \sim \frac{1}{L_{k_1}} f_{k_1}(\nu)^{L_{k_1}-1} + \frac{1}{L_{k_2} f_{k_2}(\nu)} \sum_{k\in K_*} f_k(\nu)^{L_k}, \qquad \text{as } \nu \to \infty. \tag{4}$$

The first term on the right-hand side of (4) corresponds to the asymptotic mean *escape time* $\mathbb{E}T_{(k_1,l_1),0}(\nu)$ from the initial branch $\mathcal{B}_{k_1}$, while the second term describes the contribution of the mean time spent visiting dominant branches, possibly including branch $\mathcal{B}_{k_1}$ as well. Let

$$\alpha := \lim_{\nu\to\infty} \frac{\mathbb{E}T_{(k_1,l_1),0}(\nu)}{\mathbb{E}T_{(k_1,l_1),(k_2,l_2)}(\nu)} \in [0,1] \tag{5}$$

denote the relative weight of $\mathcal{B}_{k_1}$.

Our second result gives the asymptotic distribution of the transition time $T_{(k_1,l_1),(k_2,l_2)}(\nu)$ scaled by its mean as $\nu \to \infty$.

Theorem 3.2. *If $k_1 \neq k_2$, then*

$$\frac{T_{(k_1,l_1),(k_2,l_2)}(\nu)}{\mathbb{E}T_{(k_1,l_1),(k_2,l_2)}(\nu)} \xrightarrow{d} Z, \qquad \text{as } \nu \to \infty.$$

The random variable Z can be expressed as

$$Z \stackrel{d}{=} \alpha Y + (1-\alpha)W,$$

where the random variable Y is exponentially distributed with unit mean and the random variable W is independent of Y and has a more complicated distribution, see (17), which depends on the sizes and activation rates of the dominant branches only. The possible distributions of Z are summarized in Table 1 in Section 5. In several cases the distribution of Z is exponential, which may be expected in view of the connection with many exponentiality results for the occurrence of rare events [1, 2, 3, 6, 12, 14]. In addition, we identify various cases that lead to *non*-exponentiality, typically due to the fact that the activity process spends a substantial period in branches other than k_1 and k_2.

Our third result concerns the starvation phenomenon. For $k = 1, \ldots, K$, define the random variable

$$\tau_k(t) := \int_0^t I_{\{X(s)\in\mathcal{B}_k\}} ds,$$

which measures how much time the activity process $\{X(t)\}_{t\geq 0}$ spends in branch $\mathcal{B}_k$ during the interval $[0,t]$. We can think of $\tau_k(t)$ as a measure of the *throughput* of component C_k over the time interval $[0,t]$. We speak of *complete starvation* or *zero throughput* of component C_k in $[0,t]$ when $\tau_k(t) = 0$. The next theorem provides insight into the time scales at which throughput starvation occur for a component of the network.

Theorem 3.3. *Assume $X(0) = (k_1, l_1)$ and $k_2 \neq k_1$. If $t(\nu) \sim \omega \mathbb{E}T_{(k_1,l_1),(k_2,1)}(\nu)$, with $\omega \in \mathbb{R} \cup \{0\}$, then*

$$\lim_{\nu\to\infty} \mathbb{P}\left(\tau_{k_2}(t(\nu)) = 0\right) \geq \mathbb{P}\left(Z \geq \omega\right). \tag{6}$$

In particular, if $t(\nu) = o(\mathbb{E}T_{(k_1,l_1),(k_2,1)}(\nu))$, then

$$\lim_{\nu\to\infty} \mathbb{P}\left(\tau_{k_2}(t(\nu)) = 0\right) = 1,$$

i.e. all users in C_{k_2} have zero throughput for a period of length $t(\nu)$ with probability one as $\nu \to \infty$.

The limit in (6) says that, even if there is long-term fairness among the components, for large values of ν all users in C_{k_2} will face starvation on all time scales smaller than the mean transition time from the initial component to C_{k_2}.

For the Markov process at hand, slow transitions and starvation effects are intimately related with the mixing time. The mixing time of a process is a characterization of the time required for the process to reach equilibrium. Indeed, due to the complete partite structure of the conflict graph, the process is bound to be stuck in one of the dominant branches, leading to slow convergence to equilibrium. In Section 8 we define the mixing time in terms of the total variation distance from stationarity, and prove a lower bound for a large enough activation parameter ν. This lower bound (see Proposition 8.1) indicates that the mixing time of the process is at least as large as the mean escape time from the dominant branch, which establishes a direct connection between transition times and mixing times.

Structure of the paper

The remainder of the paper is organized as follows. In Section 4 we study the activity process within a single component, where it behaves as a birth-and-death process, bringing the asymptotic behavior within the realm of classical results. In Section 5 we then leverage these results in conjunction with a geometric-sum representation to prove both Theorems 3.1 and 3.2. In Section 6 we sketch how the approach extends to scenarios where some of the users within the same component may interfere as well, relying on the same geometric-sum representation, but using more general asymptotic exponentiality results in [6] for the single-component behavior. We prove Theorem 3.3 and a complementary result for throughput ''near-saturation'' in Section 7. In Section 8 we derive a lower bound on the mixing time. Lastly, in Section 9 we make some concluding remarks and sketch some directions for further research.

4 Hitting times within a single branch

We first present a few results for the case where the two states (k_1, l_1) and (k_2, l_2) belong to the same branch, i.e. $k_1 = k_2$, and $l_1 > l_2$. In this case, the presence of the other components does not affect the transition time, and hence we focus on a single branch, dropping the component index until further notice. Within a single component of size L, the process $\{X(t)\}_{t\geq 0}$ evolves as an elementary birth-and-death process on the state space $\{L, L-1, \ldots, 1, 0\}$, so we can exploit several classical results for such processes. If we denote by $f(\nu)$ the activation rate for this component as a function of ν, then the transition rates read

$$\begin{aligned} &q(l, l+1) = (L-l)f(\nu), \quad l = 0, \ldots, L-1, \\ &q(l, l-1) = l, \quad l = 1, \ldots, L. \end{aligned}$$

4.1 Asymptotic growth rate

We first show how the mean transition time scales with the aggressiveness parameter ν.

Proposition 4.1. *For $L \geq l_1 > l_2 \geq 0$,*

$$\mathbb{E}T_{l_1,l_2}(\nu) \sim \frac{l_2!(L-l_2-1)!}{L!} f(\nu)^{L-l_2-1}, \qquad \text{as } \nu \to \infty.$$

Proof. First observe that $\mathbb{E}T_{l_1,l_2}(\nu) = \sum_{l=l_1}^{l_2+1} \mathbb{E}T_{l,l-1}(\nu)$, so we can exploit a general result for birth-and-death processes [11], which in the present case says that, for $0 < l \leq L$,

$$\mathbb{E}T_{l,l-1}(\nu) = \frac{1}{l} \sum_{n=l}^{L} \frac{\pi_n(\nu)}{\pi_l(\nu)}.$$

Now (2) implies that $\pi_n(\nu) = o(\pi_L(\nu))$ as $\nu \to \infty$ for all $n = l, \ldots, L-1$, so that

$$\mathbb{E}T_{l,l-1}(\nu) \sim \frac{1}{l} \frac{\pi_L(\nu)}{\pi_l(\nu)} = \frac{(l-1)!(L-l)!}{L!} f(\nu)^{L-l}, \qquad \text{as } \nu \to \infty.$$

Thus $\mathbb{E}T_{l,l-1}(\nu) = o(\mathbb{E}T_{l_2+1,l_2}(\nu))$ as $\nu \to \infty$ for all $l = l_1, \ldots, l_2$, and hence $\mathbb{E}T_{l_1,l_2}(\nu) \sim \mathbb{E}T_{l_2+1,l_2}(\nu)$ as $\nu \to \infty$ and the result follows. □

In order to gain insight in starvation effects, we are particularly interested in the time for the activity process to reach the center state 0, referred to as *escape time*, because at such points in time users in other components have an opportunity to activate. Proposition 4.1 shows that

$$\mathbb{E}T_{l_1,0}(\nu) \sim \frac{1}{L} f(\nu)^{L-1}, \quad \text{as } \nu \to \infty. \tag{7}$$

Hence, the mean escape time grows asymptotically as a power of $f(\nu)$, with the exponent equal to the component size minus one, and independent of the starting state l_1.

4.2 Asymptotic exponentiality

We now turn to the scaled escape time, and show that it has an asymptotically exponential distribution. We will leverage the following well-known result for birth-and-death processes, which is commonly attributed to Keilson [13] or Karlin and McGregor [10].

Theorem 4.2. *Consider a birth-and-death process with generator matrix Q on the state space $\{0,\dots,L\}$ started at state L. Assume that 0 is an absorbing state, and that the other birth rates $\{\lambda_i\}_{i=1}^{L-1}$ and death rates $\{\mu_i\}_{i=1}^{L}$ are positive. Then the absorption time in state 0 is distributed as the sum of L independent exponential random variables whose rate parameters are the L nonzero eigenvalues of $-Q$.*

Let $Q(\nu)$ be the generator matrix of the birth-and-death process $\{X(t)\}_{t\geq 0}$ on the state space $\{L, L-1,\dots,1,0\}$, with 0 an absorbing state. Let $\{\theta_i(\nu)\}_{i=1}^{L}$ denote the non-zero eigenvalues of $-Q(\nu)$. It is known [20] that these eigenvalues are distinct, real and strictly positive, so we denote $0 < \theta_1(\nu) < \theta_2(\nu) < \cdots < \theta_L(\nu)$.

Theorem 4.2 gives

$$T_{L,0}(\nu) \overset{d}{=} \sum_{i=1}^{L} Y_i(\nu), \tag{8}$$

with $Y_1(\nu),\dots,Y_L(\nu)$ independent and exponentially distributed random variables with $\mathbb{E}Y_i(\nu) = 1/\theta_i(\nu)$.

The following lemma relates the growth rates of the eigenvalues as $\nu \to \infty$ to the mean escape time $\mathbb{E}T_{L,0}(\nu)$.

Lemma 4.3. $\lim_{\nu\to\infty} \theta_i(\nu) \cdot \mathbb{E}T_{L,0}(\nu) = 1$ *if* $i = 1$ *and* ∞ *if* $i = 2,\dots,L$.

The proof of Lemma 4.3 is presented in A, and exploits detailed information about the growth rates of the eigenvalues obtained via symmetrization and the Gershgorin circle theorem. Lemma 4.3 shows that the smallest eigenvalue $\theta_1(\nu)$ becomes dominant as $\nu \to \infty$, but also proves the asymptotic exponentiality of the escape time. Indeed, denoting by $\mathcal{L}_X(s) = \mathbb{E}(e^{-sX})$, with $\mathrm{Re}(s) > 0$, the Laplace transform of a random variable X, (8) gives

$$\mathcal{L}_{T_{L,0}(\nu)/\mathbb{E}T_{L,0}(\nu)}(s) = \prod_{i=1}^{L} \Big(1 + \frac{s}{\theta_i(\nu)\cdot \mathbb{E}T_{L,0}(\nu)}\Big)^{-1}.$$

Lemma 4.3 implies that

$$\lim_{\nu\to\infty} \mathcal{L}_{T_{L,0}(\nu)/\mathbb{E}T_{L,0}(\nu)}(s) = \frac{1}{1+s}.$$

The continuity theorem for Laplace transforms then yields that the scaled escape time has an asymptotically exponential distribution as stated in the next theorem, where $\mathrm{Exp}(\lambda)$ denotes an exponentially distributed random variable with mean $1/\lambda$.

Theorem 4.4.

$$\frac{T_{L,0}(\nu)}{\mathbb{E}T_{L,0}(\nu)} \xrightarrow{d} \mathrm{Exp}(1), \quad \textit{as } \nu \to \infty.$$

This result can be understood as follows. For large ν, the probability of hitting state 0 before the first return to state L becomes small. So the time $T_{L,0}(\nu)$ consists of a geometrically distributed number of excursions from L which return to L without hitting 0, followed by the remaining part of the excursion that hits 0. Hence, apart from this final part, $T_{L,0}(\nu)$ is the sum of a large geometrically distributed number of i.i.d. random variables, which indeed is expected to be exponential.

The fact that the time until the first occurrence of a rare event is asymptotically exponential, is a widely observed phenomenon [14]. Exponentiality of the hitting time of some subset B of the state space typically arises when the probability of hitting B in a single regenerative cycle is 'small', and the cycle lengths are 'not too heavy tailed' [6, 14].

This is also true for our situation, and hence an alternative proof of Theorem 4.4 can be obtained using [6, Thm. 1] (which is a generalized version of [12]). We do not use the probabilistic approach in [6] here, because the special case of a birth-and-death process allows for explicit analysis. However, in Section 6 we will discuss how this probabilistic approach can be exploited when the individual components have a more general structure and cannot be described by birth-and-death processes.

Let us finally remark that for reversible Markov processes similar exponentiality results were established in [1]-[3]. Aldous [1] showed that a result like Theorem 4.4 can be expected when the underlying Markov process converges rapidly to stationarity. This is indeed the case for the Markov process $\{X(t)\}_{t\geq 0}$ restricted to a single branch.

To extend Theorem 4.4 to the case of a general starting state $0 < l \leq L$, we need the following technical lemma, whose proof is given in B.

Lemma 4.5. *Let $T(\nu), U(\nu), V(\nu), W(\nu)$ be non-negative random variables. Consider the properties*

(i) $\lim_{\nu\to\infty} \mathbb{E}V(\nu)/\mathbb{E}U(\nu) = \lim_{\nu\to\infty} \mathbb{E}W(\nu)/\mathbb{E}U(\nu) = 0$.

(ii) *For every $\nu > 0$, $U - V \leq_{st} T \leq_{st} U + W$, i.e. for every $t > 0$,*

$$\mathbb{P}(U - V > t) \leq \mathbb{P}(T > t) \leq \mathbb{P}(U + W > t).$$

(iii) $U(\nu)/\mathbb{E}U(\nu) \xrightarrow{d} Z$ *as $\nu \to \infty$, where Z is a continuous random variable independent of ν.*

Then,

(a) *If* (i) *and* (ii) *hold, then* $\lim_{\nu\to\infty} \mathbb{E}T(\nu)/\mathbb{E}U(\nu) = 1$.

(b) *If* (i), (ii) *and* (iii) *hold, then* $T(\nu)/\mathbb{E}T(\nu) \xrightarrow{d} Z$, *as $\nu \to \infty$.*

Proposition 4.6. *For any $0 < l \leq L$,*

$$\frac{T_{l,0}(\nu)}{\mathbb{E}T_{l,0}(\nu)} \xrightarrow{d} \mathrm{Exp}(1), \qquad \text{as } \nu \to \infty.$$

Proof. The birth-and-death structure of the process and the strong Markov property yield the stochastic identity $T_{L,0}(\nu) \stackrel{d}{=} T_{L,l}(\nu) + T_{l,0}(\nu)$, which gives the stochastic bounds $T_{L,0}(\nu) - T_{L,l}(\nu) \leq_{\mathrm{st}} T_{l,0}(\nu) \leq_{\mathrm{st}} T_{L,0}(\nu)$ (the two terms in the lower bound being dependent). It follows from Theorem 4.4 that $T_{L,0}(\nu)/\mathbb{E}T_{L,0}(\nu) \xrightarrow{d} \mathrm{Exp}(1)$ as $\nu \to \infty$. In order to complete the proof, we can then use Lemma 4.5, taking $U(\nu) = T_{L,0}(\nu)$, $V(\nu) = T_{L,l}(\nu)$ and $W(\nu) = 0$. The condition which needs to be checked is $\lim_{\nu\to\infty} \mathbb{E}V(\nu)/\mathbb{E}U(\nu) = 0$, which follows directly from Proposition 4.1. □

4.3 More general coefficients and applications

We can extend our analysis to more general activation and deactivation dynamics inside a single branch, described by

$$\begin{aligned} q(l, l+1) &= a_l f(\nu), \quad l = 1, \ldots, L-1, \\ q(l, l-1) &= d_l, \quad l = 2, \ldots, L, \end{aligned}$$

where a_l, d_l are positive real coefficients. Specifically, Proposition 4.1 can be generalized to the following result. For $L \geq l_1 > l_2 \geq 0$,

$$\mathbb{E}T_{l_1,l_2}(\nu) \sim \frac{1}{d_{l_2+1}} \Big(\prod_{i=l_2+1}^{L-1} \frac{a_i}{d_{i+1}} \Big) f(\nu)^{L-l_2-1}, \qquad \text{as } \nu \to \infty. \tag{9}$$

Also Lemma 4.3 and thus Proposition 4.6 can be shown to hold for these more general rates (see A).

These results for general coefficients have some interesting applications, beyond the model considered in this paper. One example is the continuous-time Markov process $\{M_t\}_{t\geq 0}$ on $S = \{0, 1, \ldots, c\}$, describing the number of busy servers at time t. Suppose that the service rate of each server is 1 and the arrival rate is ν which will grow large in a heavy-traffic regime. The escape time $T_{s,0}(\nu)$, choosing $a_n = 1$ and $d_n = n$, $n = 1, \ldots, c$, then describes the time it takes for this system to *drain* (i.e. to have all the servers idle) when starting with $s \geq 1$ busy servers. Then (9) gives $\mathbb{E}T_{s,0}(\nu) \sim \nu^{c-1}/c!$ as $\nu \to \infty$, which does not depend on the starting state $s \geq 1$. The scaled drain time obeys $\frac{T_{s,0}(\nu)}{\mathbb{E}T_{s,0}(\nu)} \xrightarrow{d} \mathrm{Exp}(1)$ as $\nu \to \infty$.

5 Proofs of Theorems 3.1 and 3.2

In this section we investigate the asymptotic behavior of the transition time $T_{(k_1,l_1),(k_2,l_2)}(\nu)$ as $\nu \to \infty$ for any pair of states (k_1, l_1) and (k_2, l_2), with $k_1 \neq k_2$. In Subsection 5.1 we provide a stochastic representation of the transition time, which we use to derive the asymptotic mean transition time in Subsection 5.2 leading to Theorem 3.1. In Subsection 5.3 we will obtain the asymptotic distribution of the scaled transition time leading to Theorem 3.2. In Subsection 5.4 we consider in detail the random variable W that occurs in Theorem 3.2. We give an overview of all possible forms of asymptotic behavior and the conditions under which they occur in Subsection 5.5.

5.1 Stochastic representation of the transition time

Consider the evolution of the process as it makes a transition from a state (k_1, l_1) to a state (k_2, l_2) and defining the following random variables:

- $T^{(0)}_{(k_1,1),0}(\nu)$: time to reach state 0 after state $(k_1, 1)$ is visited for the first time;
- $N_k(\nu)$: number of times the process makes a transition $0 \to (k, 1)$, $k \neq k_2$, before the first transition $0 \to (k_2, 1)$ occurs;
- $\hat{T}^{(i)}_{0,(k,1)}(\nu)$: time spent in state 0 before the i-th transition to state $(k, 1)$, $k \neq k_2$, $i = 1, \ldots, N_k(\nu)$;
- $\hat{T}_{0,(k_2,1)}(\nu)$: time spent in state 0 before the first transition to state $(k_2, 1)$;
- $T^{(i)}_{(k,1),0}(\nu)$: time to return to state 0 after the i-th transition to state $(k, 1)$, $k \neq k_2$, $i = 1, \ldots, N_k(\nu)$.

With the above definitions, it is readily seen that the following stochastic representation holds.

Proposition 5.1. *The transition time $T_{(k_1,l_1),(k_2,l_2)}$ can be represented as*

$$T_{(k_1,l_1),(k_2,l_2)} \stackrel{d}{=} T_{(k_1,l_1),(k_1,1)} + T^{(0)}_{(k_1,1),0} + \sum_{k\neq k_2} \sum_{i=1}^{N_k} \left(\hat{T}^{(i)}_{0,(k,1)} + T^{(i)}_{(k,1),0} \right) + \hat{T}_{0,(k_2,1)} + T_{(k_2,1),(k_2,l_2)}, \tag{10}$$

where the dependence on the parameter ν is suppressed for compactness and all the random variables representing time durations are mutually independent as well as independent of the random variables $N_k(\nu)$, $k \neq k_2$.

Denote $F(\nu) = \sum_{k=1}^{K} L_k f_k(\nu)$. The random variables $T^{(i)}_{(k,1),0}$ are i.i.d. copies of $T_{(k,1),0}$, $i = 1, \ldots, N_k(\nu)$, $k \neq k_2$, while the random variables $\hat{T}_{0,(k_2,1)}$ and $\hat{T}^{(i)}_{0,(k,1)}$, $k \neq k_2$, $i = 1, \ldots, N_k$, are i.i.d. copies of $T_0 \stackrel{d}{=} \mathrm{Exp}(F(\nu))$, which is the residence time in state 0. Write $X \stackrel{d}{=} \mathrm{Geo}(p)$ when X is a random variable with geometric distribution $\mathbb{P}(X = n) = p(1-p)^n$, $n \in \mathbb{N} \cup \{0\}$. Define the random variable $N(\nu) := \sum_{k\neq k_2} N_k(\nu)$, counting the total number of entrances in branches other than k_2 before hitting the target branch $\mathcal{B}_{k_2}$. For all $k = 1, \ldots, K$, denote $p_k(\nu) := L_k f_k(\nu)/F(\nu)$. Obviously,

$$N(\nu) \stackrel{d}{=} \mathrm{Geo}(p_{k_2}(\nu)),$$

while the marginal distribution of $N_k(\nu)$ is $\mathrm{Geo}(\frac{p_{k_2}(\nu)}{p_{k_2}(\nu)+p_k(\nu)})$.

We want to distinguish the branches that significantly affect the dynamics of the process (and hence the transition time) from those that do not. The quantity $\mathbb{E}N_k(\nu) \cdot \mathbb{E}T_{(k,1),0}(\nu)$, for $k \neq k_2$, is the mean time that the process spends in branch $\mathcal{B}_k$ along the transition $(0, 0) \to (k_2, l_2)$. Note that Proposition 4.1 gives

$$\mathbb{E}T_{(k,1),0}(\nu) \sim \frac{1}{L_k} f_k(\nu)^{L_k - 1}, \quad \text{as } \nu \to \infty, \tag{11}$$

and that

$$\mathbb{E}N_k(\nu) = \frac{p_k(\nu)}{p_{k_2}(\nu)} = \frac{L_k f_k(\nu)}{L_{k_2} f_{k_2}(\nu)}. \tag{12}$$

Therefore

$$\frac{\mathbb{E}N_k(\nu) \cdot \mathbb{E}T_{(k,1),0}(\nu)}{\mathbb{E}N_j(\nu) \cdot \mathbb{E}T_{(j,1),0}(\nu)} \sim \frac{f_k(\nu)^{L_k}}{f_j(\nu)^{L_j}}, \quad \text{as } \nu \to \infty,$$

which shows that indeed only the visits to dominant branches asymptotically contribute to the mean transition time.

5.2 Asymptotic mean transition time

We present here the proof of Theorem 3.1. Consider the stochastic representation (10) of the transition time $T_{(k_1,l_1),(k_2,l_2)}(\nu)$. Proposition 4.1 implies that

$$\mathbb{E}T_{(k_1,l_1),(k_1,1)}(\nu) \sim \frac{1}{L_{k_1}(L_{k_1}-1)} f_{k_1}(\nu)^{L_{k_1}-2}, \quad \text{as } \nu \to \infty,$$

and that

$$\mathbb{E}T_{(k_1,1),0}(\nu) \sim \frac{1}{L_{k_1}} f_{k_1}(\nu)^{L_{k_1}-1} \quad \text{as } \nu \to \infty.$$

Hence $\mathbb{E}T_{(k_1,l_1),(k_1,1)(\nu)} = o\left(\mathbb{E}T_{(k_1,1),0}(\nu)\right)$ as $\nu \to \infty$. Moreover $\mathbb{E}\hat{T}_{0,(k,1)}(\nu) = o(1)$. Lemma 5.2 below implies that $\mathbb{E}T_{(k_2,1),(k_2,l_2)}(\nu) = o\left(\mathbb{E}T_{(k_1,l_1),(k_2,l_2)}(\nu)\right)$. The asymptotic relation (4) then follows using the definition of K_*, the asymptotic estimate (11) and the identity (12).

The following lemma guarantees that once the process has entered the target branch $\mathcal{B}_{k_2}$, even if it may exit from it, the mean time it takes to reach the target state (k_2, l_2) is negligible with respect to the mean overall transition time.

Lemma 5.2.

$$\mathbb{E}T_{(k_2,1),(k_2,l_2)}(\nu) = o\left(\mathbb{E}T_{(k_1,l_1),(k_2,l_2)}(\nu)\right), \quad \textit{as } \nu \to \infty.$$

Proof. Consider the event

$$\begin{aligned}\mathcal{E}(\nu) &= \{\text{the first } l_2 - 1 \text{ transitions are all towards the state } (k_2, l_2)\} \\ &= \bigcap_{i=1}^{l_2-1} \{\text{the } i\text{-th transition is from } (k_2, i) \text{ to } (k_2, i+1)\}.\end{aligned}$$

Exploiting the fact that all these events are independent, we can compute

$$\begin{aligned}\mathbb{P}\left(\mathcal{E}(\nu)\right) &= \prod_{i=1}^{l_2-1} \mathbb{P}\left(\text{the } i\text{-th transition is from } (k_2, i) \text{ to } (k_2, i+1)\right) \\ &= \frac{(L_{k_2}-1)f_{k_2}(\nu)}{(L_{k_2}-1)f_{k_2}(\nu)+1} \cdot \frac{(L_{k_2}-2)f_{k_2}(\nu)}{(L_{k_2}-2)f_{k_2}(\nu)+2} \cdot \cdots \cdot \frac{(L_{k_2}-l_2+1)f_{k_2}(\nu)}{(L_{k_2}-l_2+1)f_{k_2}(\nu)+(l_2+1)},\end{aligned}$$

and clearly $\lim_{\nu\to\infty} \mathbb{P}\left(\mathcal{E}(\nu)\right) = 1$. We have that

$$\mathbb{E}\{T_{(k_2,1),(k_2,l_2)}(\nu) \mid \mathcal{E}(\nu)\} = \sum_{m=1}^{l_2-1} \frac{1}{(L_2-m)f_{k_2}(\nu)+m} =: g(\nu),$$

where $g(\nu) \downarrow 0$ as $\nu \to \infty$. Consider the events $\mathcal{E}_n^c(\nu) = \{\text{the first transition towards state 0 is the } n\text{-th one}\}$, for $n = 1, \ldots, l_2 - 1$. Note that the event $\mathcal{E}^c(\nu)$ can be decomposed as $\mathcal{E}^c(\nu) = \bigcup_{n=1}^{l_2-1} \mathcal{E}_n^c(\nu)$. Using the events $\mathcal{E}(\nu)$ and $\mathcal{E}_n^c(\nu)$, we can write

$$\mathbb{E}T_{(k_2,1),(k_2,l_2)}(\nu) = \mathbb{E}\{T_{(k_2,1),(k_2,l_2)}(\nu) \mid \mathcal{E}(\nu)\}\mathbb{P}\left(\mathcal{E}(\nu)\right) + \sum_{n=1}^{l_2-1} \mathbb{E}\{T_{(k_2,1),(k_2,l_2)}(\nu) \mid \mathcal{E}_n^c(\nu)\}\mathbb{P}\left(\mathcal{E}_n^c(\nu)\right). \tag{13}$$

Since

$$\begin{aligned}\mathbb{E}\{T_{(k_2,1),(k_2,l_2)}(\nu) \mid \mathcal{E}_n^c(\nu)\} &\le \mathbb{E}\{T_{(k_2,1),(k_2,n-1)}(\nu) \mid \mathcal{E}_n^c(\nu)\} + \mathbb{E}T_{(k_2,n-1),(k_2,l_2)}(\nu) \\ &\le \mathbb{E}\{T_{(k_2,1),(k_2,n-1)}(\nu) \mid \mathcal{E}(\nu)\} + \mathbb{E}T_{0,(k_2,l_2)}(\nu) \\ &\le \mathbb{E}\{T_{(k_2,1),(k_2,l_2)}(\nu) \mid \mathcal{E}(\nu)\} + \mathbb{E}T_{(k_1,l_1),(k_2,l_2)}(\nu),\end{aligned}$$

for $n = 1, \ldots, l_2 - 1$, it follows from (13) that

$$\mathbb{E}T_{(k_2,1),(k_2,l_2)}(\nu) \le g(\nu) + \mathbb{E}T_{(k_1,l_1),(k_2,l_2)}(\nu)\mathbb{P}\left(\mathcal{E}^c(\nu)\right).$$

We divide both sides by $\mathbb{E}T_{(k_1,l_1),(k_2,l_2)}(\nu)$, which is greater than 1 for ν sufficiently large, thanks to (7). Since $g(\nu)$ and $\mathbb{P}\left(\mathcal{E}^c(\nu)\right)$ are both $o(1)$, the proof of the lemma is complete. □

5.3 Asymptotic distribution of the transition time

We now turn to the proof of Theorem 3.2. It is clear that only the dominant branches which asymptotically contribute to the expected magnitude of the transition time will play a role, possibly along with the escape time from the initial branch. As we will show, the various dominant branches may play different roles, depending on whether the expected number of visits during the transition time is zero, $O(1)$ or infinite in the limit as $\nu \to \infty$. We introduce

$$A(\nu) := T_{(k_1,1),0}(\nu) \quad \text{and} \quad B(\nu) := \sum_{k \in K_*} \sum_{i=1}^{N_k(\nu)} T^{(i)}_{(k,1),0}(\nu), \tag{14}$$

whose means correspond to the two terms at the right-hand side of (4). From the definition (5) of the coefficient α it follows that

$$\alpha = \lim_{\nu \to \infty} \frac{\mathbb{E}A(\nu)}{\mathbb{E}A(\nu) + \mathbb{E}B(\nu)}.$$

When $\alpha = 0$ the term $A(\nu)$ becomes asymptotically negligible compared to $B(\nu)$, while the opposite holds when $\alpha = 1$. Proposition 4.6 already describes the asymptotic behavior of $A(\nu)$ after scaling. We need to understand the asymptotic behavior of $B(\nu)$ and for this purpose, it will be convenient to use a slightly different representation for it.

Define $p_*(\nu) := \sum_{k \in K_*} p_k(\nu)$ and $\hat{p}(\nu) := \frac{p_*(\nu)}{p_{k_2}(\nu)+p_*(\nu)}$. Introduce the random variable $N_*(\nu) := \text{Geo}(1-\hat{p}(\nu))$, which represents the number of visits to the dominant branches, before entering the target branch $\mathcal{B}_{k_2}$. Introduce the sequence $(\tau^{(i)}(\nu))_{i \geq 1}$ of i.i.d. random variables, $\tau^{(i)}(\nu) \stackrel{d}{=} \tau(\nu)$, where $\tau(\nu) \stackrel{d}{=} T_{(k,1),0}(\nu)$ with probability $p_k(\nu)/p_*(\nu)$ for every $k \in K_*$. Then

$$B(\nu) \stackrel{d}{=} \sum_{i=1}^{N_*(\nu)} \tau^{(i)}(\nu). \tag{15}$$

For $k \in K_*$ we define

$$\beta_k := \lim_{\nu \to \infty} \frac{L_k f_k(\nu)}{L_{k_2} f_{k_2}(\nu)}. \tag{16}$$

In view of (2), β_k may be interpreted as the stationary ratio between the number of visits to branch k and to branch k_2 as $\nu \to \infty$. Starting from branch $k_1 \neq k_2$, β_k also represents the asymptotic mean number of visits to branch $\mathcal{B}_k$ before the first entrance in branch $\mathcal{B}_{k_2}$ as $\nu \to \infty$. To avoid technicalities, we henceforth assume that all the parameters β_k are well defined. Moreover, we introduce the parameter $\beta := \sum_{k \in K_*} \beta_k$, which is the asymptotic mean number of visits to dominant branches before hitting $\mathcal{B}_{k_2}$ as $\nu \to \infty$, i.e. $\beta = \lim_{\nu \to \infty} \mathbb{E}N_*(\nu)$.

Based on the definition of the parameter β_k in (16), we partition the index set K_* of the dominant branches into three subsets, namely

$$K_* = \mathcal{N} \cup \mathcal{A} \cup \mathcal{S},$$

using the following rule:

(i) $k \in \mathcal{N}$ if $\beta_k = 0$;

(ii) $k \in \mathcal{A}$ if $\beta_k \in \mathbb{R}_+$;

(iii) $k \in \mathcal{S}$ if $\beta_k = \infty$.

The branches in $\mathcal{N}$, $\mathcal{A}$ and $\mathcal{S}$ will be called *non-attracting*, *attracting* and *strongly attracting*, respectively. Define moreover the coefficients $\gamma_{\mathcal{N}} := \sum_{k \in \mathcal{N}} \gamma_k$, $\gamma_{\mathcal{A}} := \sum_{k \in \mathcal{A}} \gamma_k$ and $\gamma_{\mathcal{S}} := \sum_{k \in \mathcal{S}} \gamma_k$, with the parameters γ_k as defined in (3).

We are now ready to present the proof of Theorem 3.2. Specifically, we prove that if $k_1 \neq k_2$, $1 \leq l_1 \leq L_{k_1}$ and $1 \leq l_2 \leq L_{k_2}$, then

$$\frac{T_{(k_1,l_1),(k_2,l_2)}(\nu)}{\mathbb{E}T_{(k_1,l_1),(k_2,l_2)}(\nu)} \stackrel{d}{\to} \alpha Y + (1-\alpha)W, \quad \text{as } \nu \to \infty,$$

where α is the constant defined in (5), Y is an exponential random variable with unit mean and W is a random variable independent of Y, with Laplace transform

$$\mathcal{L}_W(s) = \frac{1}{1 + \sum_{k \in \mathcal{A}} \frac{\gamma_k s}{1 + \gamma_k s/\beta_k} + s\gamma_{\mathcal{S}}}. \tag{17}$$

The crucial idea of the proof is to use Lemma 4.5 with the dominant term $U(\nu)$ defined as the sum of the two random variables introduced in (14), i.e. $U(\nu) := A(\nu) + B(\nu)$. Theorem 3.1 implies that $\mathbb{E}T_{(k_1,l_1),(k_2,l_2)}(\nu) \sim \mathbb{E}U(\nu)$ as $\nu \to \infty$ and its proof shows that all the other terms present in the stochastic representation (10) are negligible compared to $U(\nu)$. Note that

$$\frac{U(\nu)}{\mathbb{E}U(\nu)} = \frac{A(\nu)}{\mathbb{E}U(\nu)} + \frac{B(\nu)}{\mathbb{E}U(\nu)} = \frac{\mathbb{E}A(\nu)}{\mathbb{E}U(\nu)}\frac{A(\nu)}{\mathbb{E}A(\nu)} + \frac{\mathbb{E}B(\nu)}{\mathbb{E}U(\nu)}\frac{B(\nu)}{\mathbb{E}B(\nu)}.$$

Recall that $A(\nu)$ and $B(\nu)$ are independent by construction. If we knew that there exist two random variables Y and W such that $A(\nu)/\mathbb{E}A(\nu) \xrightarrow{d} Y$ and $B(\nu)/\mathbb{E}B(\nu) \xrightarrow{d} W$ as $\nu \to \infty$, then

$$\frac{U(\nu)}{\mathbb{E}U(\nu)} \xrightarrow{d} \alpha Y + (1-\alpha)W, \quad \text{as } \nu \to \infty,$$

and Lemma 4.5 would imply that

$$\frac{T_{(k_1,l_1),(k_2,l_2)}(\nu)}{\mathbb{E}T_{(k_1,l_1),(k_2,l_2)}(\nu)} \xrightarrow{d} \alpha Y + (1-\alpha)W, \quad \text{as } \nu \to \infty.$$

Proposition 4.6 immediately gives that

$$\frac{A(\nu)}{\mathbb{E}A(\nu)} \xrightarrow{d} Y, \quad \text{as } \nu \to \infty,$$

where Y is an exponential random variable with mean one.

Thus it remains to establish that the random variable $B(\nu)/\mathbb{E}B(\nu)$ converges to W in distribution. From the definition of $B(\nu)$ and (15), it follows that $\mathbb{E}B(\nu) = \mathbb{E}N_*(\nu)\mathbb{E}\tau(\nu)$ and that

$$\mathcal{L}_{B(\nu)/\mathbb{E}B(\nu)}(s) = G_{N_*(\nu)}\left(\mathcal{L}_{\tau(\nu)/\mathbb{E}B(\nu)}(s)\right) = G_{N_*(\nu)}\left(\mathcal{L}_{\tau(\nu)/\mathbb{E}\tau(\nu)}(s/\mathbb{E}N_*(\nu)\right), \tag{18}$$

where

$$G_{N_*(\nu)}(z) = \mathbb{E}(z^{N_*(\nu)}) = \frac{1}{1+(1-z)\mathbb{E}N_*(\nu)}. \tag{19}$$

We need to understand the asymptotic behavior of the random variable $\tau(\nu)/\mathbb{E}\tau(\nu)$. Let $T_k(\nu) = T_{(k,1),0}(\nu)$. Then $\mathcal{L}_{\tau(\nu)}(s) = \sum_{k\in K_*}\frac{p_k(\nu)}{p_*(\nu)}\mathcal{L}_{T_k(\nu)}(s)$ and hence

$$\begin{aligned}
\mathcal{L}_{\tau(\nu)/\mathbb{E}\tau(\nu)}(s/\mathbb{E}N_*(\nu)) &= \mathcal{L}_{\tau(\nu)}\Big(\frac{s}{\mathbb{E}N_*(\nu)\mathbb{E}\tau(\nu)}\Big) \\
&= \sum_{k\in K^*}\frac{p_k(\nu)}{p_*(\nu)}\mathcal{L}_{T_k(\nu)}\Big(\frac{s}{\mathbb{E}N_*(\nu)\mathbb{E}\tau(\nu)}\Big) \\
&= \sum_{k\in K^*}\frac{\mathbb{E}N_k(\nu)}{\mathbb{E}N_*(\nu)}\mathcal{L}_{T_k(\nu)/\mathbb{E}T_k(\nu)}\Big(\frac{s\mathbb{E}T_k(\nu)}{\mathbb{E}N_*(\nu)\mathbb{E}\tau(\nu)}\Big).
\end{aligned}$$

For $k \in K_*$, define

$$h_k(\nu) := \frac{\mathbb{E}T_k(\nu)}{\mathbb{E}N_*(\nu)\mathbb{E}\tau(\nu)}, \tag{20}$$

and note that $\lim_{\nu\to\infty} h_k(\nu) = \gamma_k/\beta_k$. Indeed,

$$\lim_{\nu\to\infty} h_k(\nu) = \lim_{\nu\to\infty}\frac{\mathbb{E}T_k(\nu)}{\mathbb{E}N_*(\nu)\mathbb{E}\tau(\nu)} = \lim_{\nu\to\infty}\frac{\mathbb{E}N_k(\nu)\mathbb{E}T_k(\nu)}{\mathbb{E}N_*(\nu)\mathbb{E}\tau(\nu)}\frac{1}{\mathbb{E}N_k(\nu)} = \gamma_k\left(\lim_{\nu\to\infty}\mathbb{E}N_k(\nu)\right)^{-1}.$$

Combining (18)-(20) yields

$$\begin{aligned}
\mathcal{L}_{B(\nu)/\mathbb{E}B(\nu)}(s) &= \Big[1+\Big(1-\mathcal{L}_{\tau(\nu)/\mathbb{E}\tau(\nu)}(s/\mathbb{E}N_*(\nu))\Big)\mathbb{E}N_*(\nu)\Big]^{-1} \\
&= \left[1+\Big(1-\sum_{k\in K^*}\frac{\mathbb{E}N_k(\nu)}{\mathbb{E}N_*(\nu)}\mathcal{L}_{T_k(\nu)/\mathbb{E}T_k(\nu)}\Big(\frac{s\mathbb{E}T_k(\nu)}{\mathbb{E}N_*(\nu)\mathbb{E}\tau(\nu)}\Big)\Big)\mathbb{E}N_*(\nu)\right]^{-1} \\
&= \left[1+\Big(\mathbb{E}N_*(\nu)-\sum_{k\in K^*}\mathbb{E}N_k(\nu)\mathcal{L}_{T_k(\nu)/\mathbb{E}T_k(\nu)}(sh_k(\nu))\Big)\right]^{-1} \\
&= \left[1+\sum_{k\in K^*}\mathbb{E}N_k(\nu)\left(1-\mathcal{L}_{T_k(\nu)/\mathbb{E}T_k(\nu)}(sh_k(\nu))\right)\right]^{-1}.
\end{aligned} \tag{21}$$

In order to obtain an explicit expression for $\mathcal{L}_{B(\nu)/\mathbb{E}B(\nu)}(s)$ as $\nu \to \infty$, we need the following technical lemma, which is proved in C.

Lemma 5.3. (a) *If $k \in \mathcal{S}$, then*

$$\lim_{\nu\to\infty} \mathbb{E}N_k(\nu)\left(1 - \mathcal{L}_{T_k(\nu)/\mathbb{E}T_k(\nu)}(sh_k(\nu))\right) = \gamma_k s.$$

(b) *If $k \in \mathcal{A}$, then*

$$\lim_{\nu\to\infty} \mathbb{E}N_k(\nu)\left(1 - \mathcal{L}_{T_k(\nu)/\mathbb{E}T_k(\nu)}(sh_k(\nu))\right) = \frac{\gamma_k s}{1+\gamma_k s/\beta_k}.$$

(c) *If $k \in \mathcal{N}$, then*

$$\lim_{\nu\to\infty} \mathbb{E}N_k(\nu)\left(1 - \mathcal{L}_{T_k(\nu)/\mathbb{E}T_k(\nu)}(sh_k(\nu))\right) = 0.$$

From Lemma 5.3 and (21) it follows that

$$\mathcal{L}_W(s) = \lim_{\nu\to\infty} \mathcal{L}_{B(\nu)/\mathbb{E}B(\nu)}(s) = \left[1 + \sum_{k\in\mathcal{A}} \frac{\gamma_k s}{1+\gamma_k s/\beta_k} + \sum_{k\in\mathcal{S}} \gamma_k s\right]^{-1}.$$

The independence of Y and W easily follows from the independence of the corresponding terms in the stochastic representation (10).

5.4 The random variable W: properties and interpretation

The random variable W is defined by its Laplace transform, see (17). We remark that the shape of the distribution W is fully determined by the branches in $\mathcal{A}$ and $\mathcal{S}$, independently of the branches in $\mathcal{N}$. Indeed the random variable W can be represented as

$$W \stackrel{d}{=} (1-\gamma_{\mathcal{N}})\overline{W},$$

where $\overline{W}$ is a unit-mean random variable that in no way depends on the parameters of the branches in the set $\mathcal{N}$. On the other hand, the presence of the factor $(1-\gamma_{\mathcal{N}})$ reflects the fact that the branches in $\mathcal{N}$ do affect the mean of the asymptotic scaled transition time: indeed convergence of the first moments holds if and only if $\alpha = 1$ or $\mathcal{N} = \emptyset$. Indeed,

$$\alpha\mathbb{E}Y + (1-\alpha)\mathbb{E}W = \alpha + (1-\alpha)(1-\gamma_{\mathcal{N}}),$$

and, if $\mathcal{N} \neq \emptyset$, then $\gamma_{\mathcal{N}} > 0$ and so $\alpha\,\mathbb{E}Y + (1-\alpha)\,\mathbb{E}W < 1$ when $\alpha \neq 1$.

Whenever either $\mathcal{A}$ or $\mathcal{S}$ is empty, the distribution of W is known explicitly, cf. Table 1. However, also in the scenario where both $\mathcal{A}$ and $\mathcal{S}$ are non-empty, it is still possible to give an interpretation of the distribution of W.

If $\mathcal{A} \neq \emptyset$, define $m := |\mathcal{A}|$ and label the branches belonging to $\mathcal{A}$ as $1, 2, \ldots, m$. Let $\beta_{\mathcal{A}} := \sum_{k=1}^m \beta_k \in (0,\infty)$ be the asymptotic mean number of visits to attracting branches as $\nu \to \infty$. Consider a hyper-exponentially distributed random variable H with rates β_k/γ_k and probabilities $\beta_k/\beta_{\mathcal{A}}$, $k = 1, \ldots, m$, whose Laplace transform is

$$\mathcal{L}_H(s) = \sum_{k=1}^m \frac{\beta_k}{\beta_{\mathcal{A}}} \frac{\beta_k/\gamma_k}{\beta_k/\gamma_k + s}.$$

Furthermore consider a marked Poisson process with rate $\lambda = \beta_{\mathcal{A}}/\gamma_{\mathcal{S}}$ and i.i.d. marks distributed according to H. The random variable W in Equation (17) corresponds to the sum of a random time $\mathcal{T}$, with $\mathcal{T}$ exponentially distributed with mean $1/\mu = \gamma_{\mathcal{S}}$, and the total size $\mathcal{W}(\mathcal{T})$ of the marks associated with all the events in interval $[0, \mathcal{T}]$. Indeed

$$\begin{aligned}\mathcal{L}_{\mathcal{T}+\mathcal{W}(\mathcal{T})}(s) &= \int_{t=0}^{\infty} e^{-st} e^{\lambda t\left(\sum_{k=1}^m \frac{\beta_k}{\beta_{\mathcal{A}}}\frac{\beta_k/\gamma_k}{\beta_k/\gamma_k+s}-1\right)} \mu e^{-\mu t} dt = \left[1 + \frac{\lambda}{\mu}\left(\sum_{k=1}^m \frac{\beta_k}{\beta_{\mathcal{A}}}\frac{\beta_k/\gamma_k}{\beta_k/\gamma_k+s}\right) + \frac{s}{\mu}\right]^{-1} \\ &= \left[1 + \beta_{\mathcal{A}}\left(\sum_{k=1}^m \frac{\beta_k}{\beta_{\mathcal{A}}}\frac{s}{\beta_k/\gamma_k+s}\right) + s\gamma_{\mathcal{S}}\right]^{-1} = \left[1 + \sum_{k\in\mathcal{A}} \frac{\gamma_k s}{1+\gamma_k s/\beta_k} + s\gamma_{\mathcal{S}}\right]^{-1}.\end{aligned}$$

The stochastic equality $W = \mathcal{T} + \mathcal{W}(\mathcal{T})$ may be interpreted as follows. Define $p_{\mathcal{A}} := \sum_{k\in\mathcal{A}} p_k(\nu)$ and $p_{\mathcal{S}} := \sum_{k\in\mathcal{S}} p_k(\nu)$. The total number of visits during the transition time to the branches in $\mathcal{S}$ is geometrically distributed

with parameter $p_{k_2}/p_{\mathcal{S}}$. Since the durations of these visits are independent and each relatively short compared to the transition time, the total normalized amount of time spent in the branches in $\mathcal{S}$ is exponentially distributed in the limit as $\nu \to \infty$ with mean $\gamma_{\mathcal{S}}$. The visits to the branches in $\mathcal{S}$ are interspersed with visits to the branches in $\mathcal{A}$. The number of visits to branches in $\mathcal{S}$ between two consecutive visits to branches in $\mathcal{A}$ is geometrically distributed with parameter $p_{\mathcal{A}}/p_{\mathcal{S}}$. The normalized durations of the visits to the branches in $\mathcal{A}$ have the hyper-exponential distribution H as specified above. By similar arguments as mentioned above, the normalized amounts of time between these visits are independent and exponentially distributed in the limit as $\nu \to \infty$ with mean $\gamma_{\mathcal{S}} \cdot p_{k_2}/p_{\mathcal{A}} = \gamma_{\mathcal{S}}/\beta_{\mathcal{A}}$. In other words, the visits to the branches in $\mathcal{A}$ occur as a Poisson process with rate $\lambda = \beta_{\mathcal{A}}/\gamma_{\mathcal{S}}$.

5.5 An overview of the possible limiting distributions

In this subsection we present an overview of all the possible limiting distributions of the scaled transition time by means of Table 1 and some simulation results to illustrate our findings.

Denote by $(H_i)_{i\in\mathbb{N}}$ a sequence of i.i.d. hyper-exponential random variables, $H_i \overset{d}{=} H$, while $\mathcal{G}$ is a geometric random variable $\mathcal{G} \overset{d}{=} \text{Geo}\big(\frac{1}{1+\beta_{\mathcal{A}}}\big)$, independent of all the other random variables.

α	$\mathcal{A}$	$\mathcal{S}$	Limiting distribution	Scenario
0	$\emptyset$	$\emptyset$	δ_0 (trivial r.v. identical to 0)	1a
	non-empty	$\emptyset$	$\sum_{i=1}^{\mathcal{G}} H_i(p_1,\ldots,p_m,\lambda_1,\ldots,\lambda_m)$	1b
			$\sum_{i=1}^{\mathcal{G}} \text{Exp}_i(\lambda)$ if $\beta_k/\gamma_k = \lambda \quad \forall\, k \in \mathcal{A}$	1b*
	$\emptyset$	non-empty	$\text{Exp}(1/\gamma_{\mathcal{S}})$	1c
	non-empty	non-empty	W	1d
$(0,1)$	$\emptyset$	$\emptyset$	$\text{Exp}(1/\alpha)$	2a
	non-empty	$\emptyset$	$\text{Exp}(1/\alpha) + \sum_{i=1}^{\mathcal{G}} H_i\big(\frac{\beta_1}{\beta_{\mathcal{A}}},\ldots,\frac{\beta_m}{\beta_{\mathcal{A}}},\frac{\beta_1}{(1-\alpha)\gamma_1},\ldots,\frac{\beta_m}{(1-\alpha)\gamma_m}\big)$	2b
			$\text{Exp}(1/\alpha) + \sum_{i=1}^{\mathcal{G}} \text{Exp}_i\big(\frac{\lambda}{1-\alpha}\big)$ if $\beta_k/\gamma_k = \lambda \quad \forall\, k \in \mathcal{A}$	2b*
			$\text{Exp}\big(\frac{1}{\alpha(1+\sum_{k=1}^m \beta_k)}\big)$ if $\beta_k/\gamma_k = \frac{1-\alpha}{\alpha} \quad \forall\, k \in \mathcal{A}$	2b**
			$\text{Exp}(1)$ if $\beta_k/\gamma_k = \frac{1-\alpha}{\alpha} = \sum_{i=1}^m \beta_k \quad \forall\, k \in \mathcal{A}$	2b***
	$\emptyset$	non-empty	$\text{Exp}(1/\alpha) + \text{Exp}(1/(1-\alpha)\gamma_{\mathcal{S}})$	2c
			$\text{Erlang}(2, 1/\alpha)$ if $\alpha = \gamma_{\mathcal{S}}/(1+\gamma_{\mathcal{S}})$	2c*
	non-empty	non-empty	$\text{Exp}(1/\alpha) + (1-\alpha)W$	2d
1	-	-	$\text{Exp}(1)$	3

Table 1: Overview of the possible asymptotic distributions of the scaled transition time.

The case $\alpha = 1$ always yields asymptotic exponentiality: this happens when the escape time from branch $\mathcal{B}_{k_1}$ dominates the total transition time. As soon as $\alpha \neq 1$, the set of dominant branches starts to play an important role. In particular, the shape of the asymptotic distribution depends only on the branches in the sets $\mathcal{A}$ and $\mathcal{S}$ and changes substantially whenever one of these two subsets (or both) are empty.

In the case $\alpha = 0$ a diverse range of behaviors may occur, with asymptotic exponentiality only in a somewhat degenerate special case 1c. The behavior for $\alpha \in (0,1)$ is just a weighted combination of the extreme cases $\alpha = 0$ and $\alpha = 1$, as described in Theorem 3.2. It does not give rise to fundamentally different behavior, but interestingly enough, it does yield asymptotic exponentiality in some very special cases.

If all users have the same activation rate, no matter which component they belong to, then without loss of generality, we may assume $f_k(\nu) = \nu$, $k = 1,\ldots,K$. Under this homogeneity assumption, the sizes of components

become crucial. Indeed, if one defines $L_* := \max_{k \neq k_2} L_k$ to be the size of the largest component, then $K_* = \{k \neq k_2 : L_k = L_*\}$. In this case the orders of magnitude of the two dominant terms of the stochastic representation (10) are

$$\mathbb{E}A(\nu) \sim \frac{\nu^{L_{k_1}-1}}{L_{k_1}} \quad \text{and} \quad \mathbb{E}B(\nu) \sim \frac{|K_*|\nu^{L_*-1}}{L_{k_2}}, \quad \text{as } \nu \to \infty,$$

and hence for $1 \leq l_1 \leq L_{k_1}$, $1 \leq l_2 \leq L_{k_2}$ and $k_1 \neq k_2$,

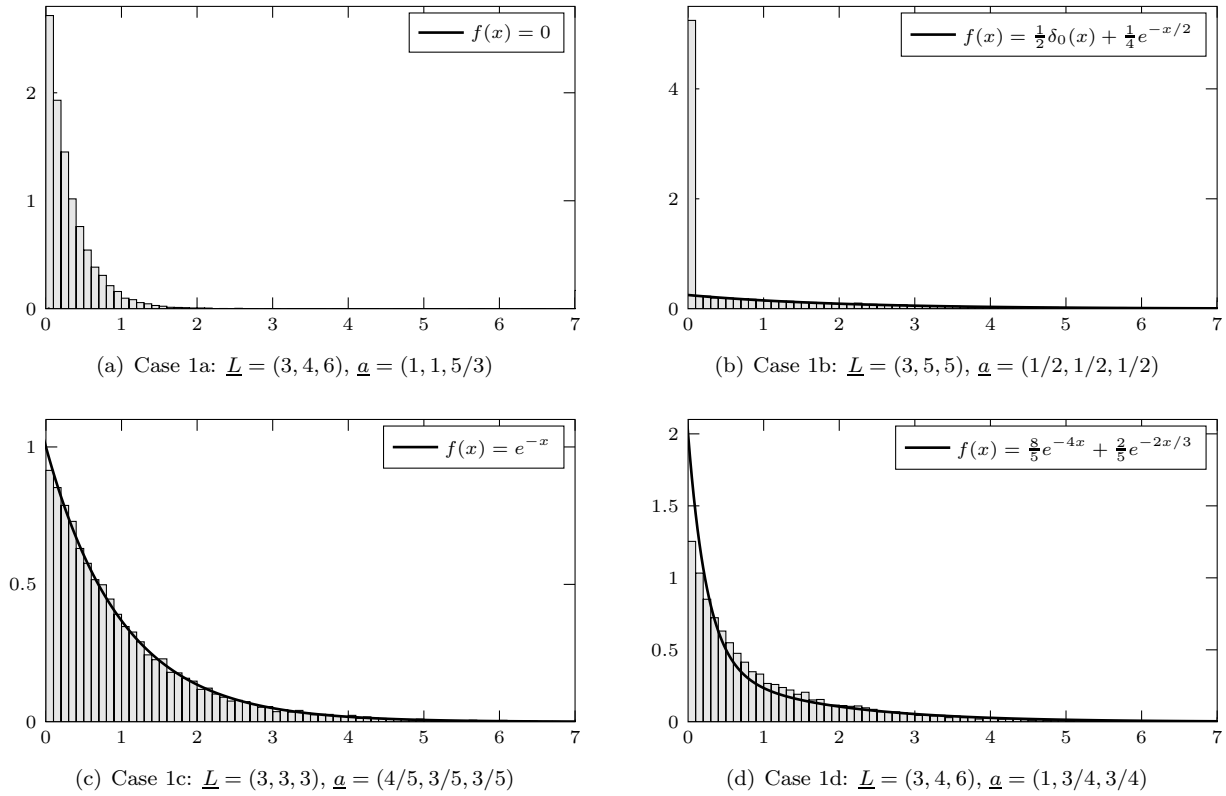
$$\mathbb{E}T_{(k_1,l_1),(k_2,l_2)}(\nu) \sim \left(\frac{I_{\{k_1 \in K_*\}}}{L_*} + \frac{|K_*|}{L_{k_2}}\right)\nu^{L_*-1}, \quad \text{as } \nu \to \infty.$$

Moreover, $\beta_k/\gamma_k = (1-\alpha)/\alpha$ for every $k \in \mathcal{A}$ and thus only two possible scenarios can occur, namely 1b* and 2b*** (see [28]). The discriminating factor between these two scenarios is α. If $k_1 \notin K_*$, then $\alpha = 0$ and thus we are in scenario 1b*. If instead $k_1 \in K_*$, then $\alpha = L_{k_2}/(|K_*|L_*)$, which means that scenario 2b*** occurs and hence asymptotic exponentiality arises.

To illustrate the range of possible limiting distributions, we present some simulation results. We consider the simplest system that is sufficiently rich to show the wide range of behaviors presented in Table 1. Specifically, we consider a complete 3-partite network, whose three components have sizes L_1, L_2 and L_3, and assume that the process starts in state $(1, L_1)$ and the target state is $(3, L_3)$.

We use activation rates of the form $f_k(\nu) = \nu^{a_k}$. For compactness, we write $\underline{a}$ for (a_1, a_2, a_3) and $\underline{L}$ for (L_1, L_2, L_3). This choice for the activation rates allows to invert the Laplace transform of W in all the cases and thus obtain a probability density function $f(x)$, which can be compared with the simulation data.

All the simulations have been performed choosing the parameter $\nu = 150$ and simulating the transition time for each network 20000 times with a customized code written in the programming language C. The results are shown in Figures 3 and 4.



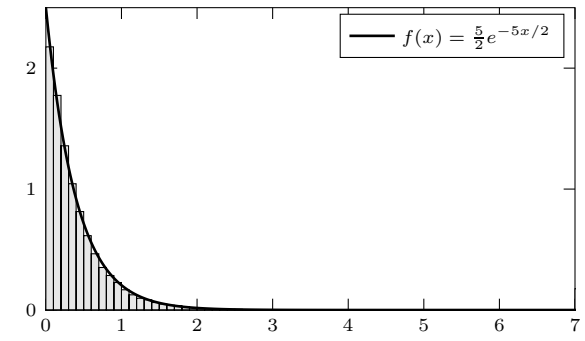
(a) Case 1a: $\underline{L} = (3, 4, 6)$, $\underline{a} = (1, 1, 5/3)$

(b) Case 1b: $\underline{L} = (3, 5, 5)$, $\underline{a} = (1/2, 1/2, 1/2)$

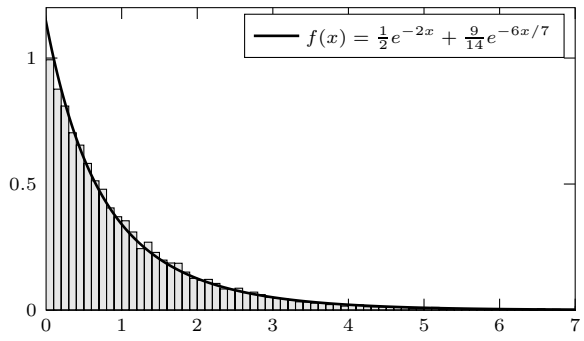
(c) Case 1c: $\underline{L} = (3, 3, 3)$, $\underline{a} = (4/5, 3/5, 3/5)$

(d) Case 1d: $\underline{L} = (3, 4, 6)$, $\underline{a} = (1, 3/4, 3/4)$

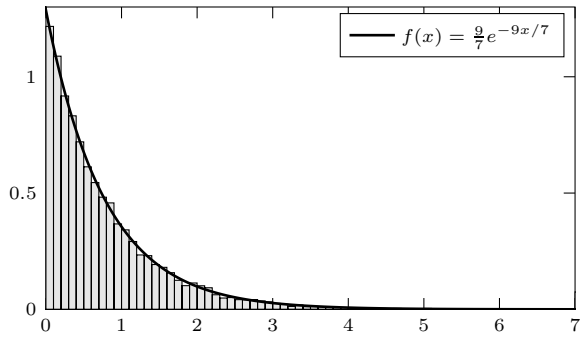
Figure 3: Plots of the empirical probability density function of the scaled transition times and the density $f(x)$ of $\alpha Y + (1-\alpha)W$.



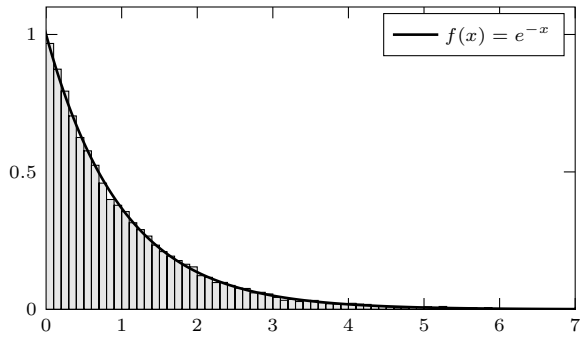
(a) Case 2a: $\underline{L} = (3, 4, 2)$, $\underline{a} = (9/10, 9/10, 9/5)$



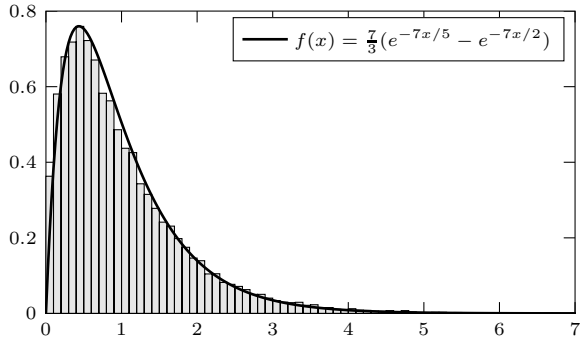
(b) Case 2b*: $\underline{L} = (4, 3, 4)$, $\underline{a} = (2/3, 1, 1)$



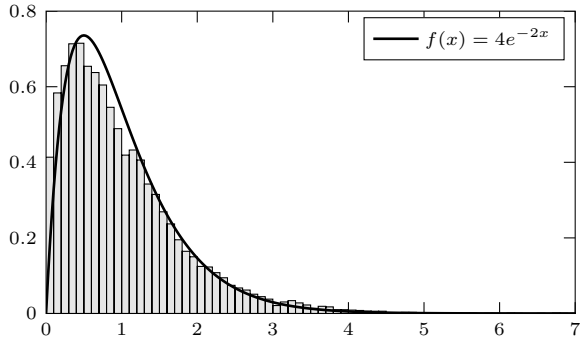
(c) Case 2b**: $\underline{L} = (2, 4, 5)$, $\underline{a} = (7/4, 7/8, 7/4)$



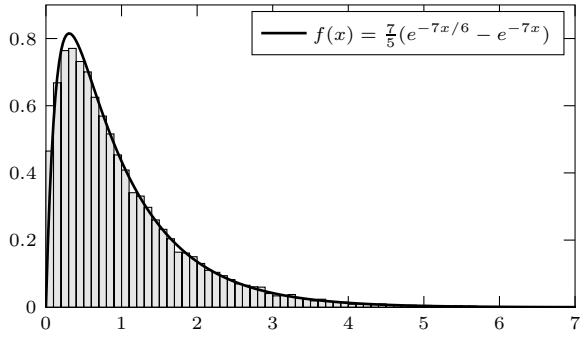
(d) Case 2b***: $\underline{L} = (3, 3, 5)$, $\underline{a} = (7/8, 7/8, 7/8)$



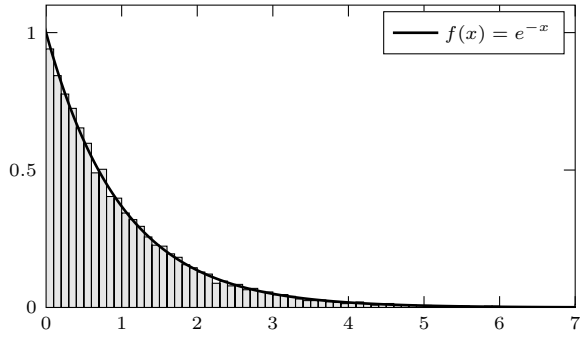
(e) Case 2c: $\underline{L} = (5, 2, 2)$, $\underline{a} = (4/9, 4/3, 8/9)$



(f) Case 2c*: $\underline{L} = (5, 2, 5)$, $\underline{a} = (1/2, 3/2, 1)$



(g) Case 2d: $\underline{L} = (4, 2, 6)$, $\underline{a} = (3/5, 6/5, 3/5)$



(h) Case 3: $\underline{L} = (3, 3, 3)$, $\underline{a} = (1, 3/4, 3/2)$

Figure 4: Plots of the empirical probability density function of the scaled transition times and the density $f(x)$ of $\alpha Y + (1 - \alpha)W$.

6 Model extensions

So far we have assumed that two users interfere if and only if they belong to different components. In this section, we continue to assume that users that belong to different components interfere, but we allow users within the same component to interfere as well. If two or more users within component C_k interfere with each other, there will be fewer admissible activity configurations of smaller size. In particular, not all the L_k users of component C_k can no longer be active simultaneously, which would then ease the transitions among different components.

In the previous sections, we further assumed all users within the same component to have the same activation rate, so that state aggregation could be applied to obtain an equivalent Markov process with a star-shaped state space. In this section, we allow the users within the same component to have possibly different activation rates. With minor abuse of notation, denote by $f_l(\nu)$ the activation rate of user l, and define $F_k(\nu) := \sum_{l \in C_k} f_l(\nu)$ as the aggregate activation rate of all users in the k-th component.

The components are assumed to be *minimal*, in the sense that they cannot be split into two non-trivial components, while retaining the full interference across components. As before, each independent set of the conflict graph must be a subset of one of the components, because two users that belong to different components by definition interfere. However, some subsets within the same component may no longer be independent sets in the conflict graph.

For every $x \in \Omega^*$ define $V_x \subseteq V$ to be the subset of users which are active in configuration x, i.e. $V_x := \{i \in V : x_i = 1\}$. For every $x \in \Omega^*$, V_x is by construction an independent set in the conflict graph G. Define moreover $\Omega_k := \{x \in \Omega^* : V_x \subseteq C_k, x \neq 0\}$. Then

$$\Omega^* = \{0\} \cup \bigcup_{k \in K} \Omega_k.$$

By construction, it follows that each component satisfies a certain monotonicity property: If $V_x \in \Omega_k$, then for every nonempty $V_y \subseteq V_x$, we have $V_y \in \Omega_k$. Indeed, if x a feasible configuration and it belongs to Ω_k, then any configuration obtained from it by switching off some users is still feasible and belongs to Ω_k as well. The next lemma shows that between any pair of activity states corresponding to the same component, there continues to exist a path which does not visit the state $0 \in \Omega^*$.

Lemma 6.1. *If $V_x, V_y \subseteq C_k$, $V_x, V_y \neq \emptyset$, then there exists a sequence of transitions between x and y that does not pass through the state $0 \in \Omega^*$.*

Proof. If $V_x \cap V_y \neq \emptyset$, then the statement is trivially true. Suppose instead that $V_x \cap V_y = \emptyset$ and without loss of generality take $V_x = \{a_1, \dots, a_l\}$ and $V_y = \{b_1, \dots, b_m\}$. Define moreover $\mathcal{C} := C_k \setminus (V_x \cup V_y) = \{c_1, \dots, c_n\}$. Thanks to the monotonicity property of C_k, we have that $V_{a_1}, \dots, V_{a_l}, V_{b_1} \dots, V_{b_m}$ all belong to C_k and each of them can be reached from V_x and V_y, respectively, without passing through the state $0 \in \Omega^*$. Suppose that every path from x to y passes through state $0 \in \Omega^*$. Then none of the configurations $V_{\{a_k, b_j\}}$, $k = 1, \dots, l$, $j = 1, \dots, m$, belongs to C_k, i.e. every user in $\{a_1, \dots, a_l\}$ interferes with every user in $\{b_1, \dots, b_m\}$.

We claim that every $c \in \mathcal{C}$ interferes either with all the users in $\{a_1, \dots, a_l\}$ or with all the users in $\{b_1, \dots, b_m\}$. Indeed, if there exist $a \in V_x$, $b \in V_y$ such that both a and b do not interfere with c, then we could construct a path from x to y which does not pass through state 0, namely $V_x \to \cdots \to \{a\} \to \{a, c\} \to \{c\} \to \{b, c\} \to \{b\} \to \cdots \to V_y$. If there exists a user $c \in \mathcal{C}$ that interferes with all the users in $\{a_1, \dots, a_l\}$ *and* $\{b_1, \dots, b_m\}$, then the component C_k would not be minimal, since it can be split in $C_k \setminus \{c\}$ and $\{c\}$. Thus every $c \in \mathcal{C}$ interferes either with all the users in $\{a_1, \dots, a_l\}$ or with all the users in $\{b_1, \dots, b_m\}$, but not both. We can then consider two sets $A = V_x \cup \mathcal{C}_A$ and $B = V_y \cup \mathcal{C}_B$, with $\mathcal{C}_A \cap \mathcal{C}_B = \emptyset$ and $\mathcal{C}_A \cup \mathcal{C}_B = \mathcal{C}$, such that users in A interfere with all users in B and vice versa. Therefore $C_k = A \cup B$ is not a minimal component. □

Clearly state $0 \in \Omega^*$ continues to be a bottleneck state which must be visited along any path between activity states corresponding to different components. For a user $l \in V$, denote by e_l the configuration in Ω where only the user l is active. Clearly, $l \in C_k$ if and only if $e_l \in \Omega_k$. For any two states $x, y \in \Omega^*$, let $T_{x,y} = \inf\{t \geq 0 : X(t) = y | X(0) = x\}$ be a random variable representing the transition time from state x to state y, i.e. the hitting time of state y starting in state x.

Let $x, y \in \Omega^*$ be two activity states, with $V_x \subseteq C_{k_1}$ and $V_y \subseteq C_{k_2}$, $k_1 \neq k_2$. In order to give a stochastic representation of the transition time $T_{x,y}$, similar in spirit to (10), we define the following random variables for $l \in V \setminus (\{0\} \cup C_{k_2})$:

- $N_l(\nu)$: number of times the process makes a transition $0 \to e_l \in \Omega_k$, $k \neq k_2$, before the first transition to C_{k_2} occurs;
- $\hat{T}^{(i)}_{0,e_l}(\nu)$: time spent in state 0 before the i-th transition to state $e_l \in \Omega_k$, with $k \neq k_2$, $i = 1, \dots, N_l(\nu)$;

- $T^{(i)}_{e_l,0}(\nu)$: time to return to state 0 after the i-th transition to state $e_l \in \Omega_k$, with $k \neq k_2$, $i = 1, \ldots, N_l(\nu)$.

Moreover, for $l \in C_{k_2}$, define $\hat{T}_{0,e_l}(\nu)$ as the time spent in state 0 before the first transition to state $e_l \in \Omega_{k_2}$. Lemma 6.1 implies that the transition time $T_{x,y}$ may be represented as

$$T_{x,y} = T_{x,0} + \sum_{k \neq k_2} \sum_{l \in C_k} \sum_{i=1}^{N_l} (\hat{T}^{(i)}_{0,e_l} + T^{(i)}_{e_l,0}) + \sum_{l \in C_{k_2}} I_l (\hat{T}_{0,e_l} + T_{e_l,y}), \tag{22}$$

where I_l, $l \in C_{k_2}$, are 0-1 variables with $\sum_{l \in C_{k_2}} I_l = 1$ and $\mathbb{P}(I_l = 1) = f_l(\nu)/F_{k_2}(\nu)$, $l \in C_{k_2}$, and the variables $T_{e_l,0}$, $l \in C_k$, are transition times when considering the component C_k in isolation. Moreover, in the above representation the dependence on the parameter ν is suppressed for compactness and all the random variables representing time durations are mutually independent as well as independent of the random variables $N_l(\nu)$, $l \in V \setminus (\{0\} \cup C_{k_2})$.

In order to determine the asymptotic behavior of the transition time $T_{x,y}(\nu)$ as $\nu \to \infty$, we now proceed to analyze the asymptotic behavior of the escape times $T_{e_l,0}$, $l \in C_k$, in the stochastic representation. Unless stated otherwise, we henceforth let $z \in \Omega_k$ and focus on the Markov process $\{X^*(t)\}_{t \geq 0}$ restricted to the state space $\Omega_k^+ = \Omega_k \cup \{0\}$.

The steady-state probability of a state $u \in \Omega_k^+$ is

$$\pi_u(\nu) = \frac{1}{Z_k(\nu)} \prod_{l \in C_k} f_l(\nu)^{u_l},$$

with normalization constant

$$Z_k(\nu) = \sum_{u' \in \Omega_k^+} \prod_{l \in C_k} f_l(\nu)^{u'_l}.$$

Define

$$g_k(\nu) := \max_{u \in \Omega_k} \prod_{l \in C_k} f_l(\nu)^{u_l} \quad \text{and} \quad \eta_k := \min_{l \in C_k} \lim_{\nu \to \infty} \log_\nu f_l(\nu).$$

We make the mild technical assumptions that $\eta \in (0, \infty)$ and that $\psi_k = \lim_{\nu \to \infty} Z_k(\nu)/g_k(\nu)$ exists. Then the following two asymptotic properties of the escape time $T_{z,0}(\nu)$ can be established:

$$\mathbb{E}T_{z,0}(\nu) \sim \frac{\psi_k g_k(\nu)}{F_k(\nu)}, \quad \text{as } \nu \to \infty, \tag{23}$$

and

$$\frac{T_{z,0}(\nu)}{\mathbb{E}T_{z,0}(\nu)} \xrightarrow{d} \text{Exp}(1), \quad \text{as } \nu \to \infty. \tag{24}$$

In order to provide a brief sketch of the proof arguments, we first introduce some further useful notation. Let $N_{z,0}(\nu)$ be a random variable representing the number of visits to state 0 in between two consecutive visits to state z. Let $R_z(\nu)$ be the residence time in state z and $T^+_{z,z}(\nu)$ the first return time to state z. Noting that

$$\pi_z(\nu) = \frac{\mathbb{E}R_z(\nu)}{\mathbb{E}T^+_{z,z}(\nu)}, \quad \mathbb{E}R_z(\nu) \leq 1 \quad \text{and} \quad \pi_z(\nu) = \frac{1}{Z_k(\nu)} \prod_{l \in C_k} f_l(\nu)^{z_l},$$

we obtain

$$\frac{\nu^\eta \mathbb{E}T^+_{z,z}(\nu)}{\psi_k g_k(\nu)} = \frac{\nu^\eta \mathbb{E}R_z(\nu)}{\pi_z(\nu)\psi_k g_k(\nu)} \leq \frac{\nu^\eta}{\prod_{l \in C_k} f_l(\nu)^{z_l}} \frac{Z_k(\nu)}{\psi_k g_k(\nu)} = o(1), \quad \text{as } \nu \to \infty.$$

Using similar arguments as in [22], it may be shown that

$$\mathbb{P}\left(T_{u,z}(\nu) > g_k(\nu)\nu^{-\eta_k/2}\right) \leq r < 1$$

for all states u with $V_u \in \mathcal{B}_k$, implying (by the strong Markov property)

$$\mathbb{P}\left(T_{u,z}(\nu) > n g_k(\nu)\nu^{-\eta_k/2}\right) \leq r^n,$$

and that $\mathbb{P}\left(T_{0,z}(\nu) < T^+_{0,0}(\nu)\right) \geq s(\nu)$, with $s(\nu) \to 1$ as $\nu \to \infty$. This means that after a visit to state 0, the number of additional visits to that state before the first visit to state z is stochastically bounded from above by a geometrically distributed random variable with parameter $1 - s(\nu)$. This implies

$$\mathbb{E}N_{z,0}(\nu) \sim \mathbb{P}(N_{z,0}(\nu) \geq 1) \quad \text{as } \nu \to \infty. \tag{25}$$

It may then be deduced that the distribution of $T^+_{z,z}(\nu)$ satisfies the uniform integrability condition in [6]. Theorem 1 in [6] then yields the asymptotic exponentiality property in (24) and

$$\mathbb{E}T_{z,0}(\nu) \sim \frac{\mathbb{E}T^+_{z,z}(\nu)}{\mathbb{P}\left(N_{z,0}(\nu) \geq 1\right)}.$$

Observing that

$$\frac{\pi_0(\nu)}{\mathbb{E}R_0(\nu)} = \mathbb{E}N_{z,0}(\nu)\frac{\pi_z(\nu)}{\mathbb{E}R_z(\nu)},$$

and invoking (25), we deduce that the term in the right-hand side asymptotically behaves as

$$\frac{\mathbb{E}T^+_{z,z}(\nu)\pi_z(\nu)\mathbb{E}R_0(\nu)}{\pi_0(\nu)\mathbb{E}R_z(\nu)} = \frac{\mathbb{E}R_0(\nu)}{\pi_0(\nu)} = \frac{Z_k(\nu)}{F_k(\nu)} \sim \frac{\psi_k g_k(\nu)}{F_k(\nu)},$$

yielding (23) as stated.

The two asymptotic properties (23), (24) for the order-of-magnitude and the scaled distribution of the escape time $T_{z,0}(\nu)$ mirror those stated in (7) and Proposition 4.6. Using these two properties and the stochastic representation (22), similar results can be established for the asymptotic behavior of the transition time $T_{x,y}(\nu)$ as in Theorems 3.1 and 3.2. For any $l \in C_k$, define

$$\Theta_l(\nu) = \frac{f_l(\nu)g_k(\nu)}{F_k(\nu)}.$$

In this case the set K^* needs to be defined as those $l \in \bigcup_{k\neq k_2} C_k$ such that $\lim_{\nu\to\infty} \Theta_l(\nu)/\Theta_m(\nu) > 0$ for all $m \in \bigcup_{k\neq k_2} C_k$. Also, additional conditions need to be imposed in order to ensure that

$$\sum_{l\in C_{k_2}} \frac{f_l(\nu)}{F_{k_2}(\nu)}\mathbb{E}T_{e_l,y}(\nu) = o(\mathbb{E}T_{x,y}(\nu)), \qquad \text{as } \nu \to \infty,$$

which guarantees that the expected time to reach state y, once the process hits the target component C_{k_2}, is asymptotically small with respect to the overall transition time $T_{x,y}(\nu)$.

7 Throughput starvation and near-saturation

In this section we return to a complete partite networks with equal activation rates for users in the same component. We show how the results for the asymptotics of the transition time $T_{(k_1,l_1),(k_2,l_2)}(\nu)$ in Theorems 3.1 and 3.2 can be exploited to gain insight about phenomena like throughput starvation or near-saturation. More specifically, in Subsection 7.1 we present the proof of Theorem 3.3, which gives an asymptotic lower bound on the probability of throughput starvation, while in Subsection 7.2 we prove Proposition 7.3, a complementary result which indicates over what time scales throughput near-saturation occurs.

7.1 Proof of Theorem 3.3

Observe that $\tau_{k_2}(t(\nu)) > 0$ implies $t(\nu) > T_{(k_1,l_1),(k_2,1)}(\nu)$, because the throughput of branch $\mathcal{B}_{k_2}$ remains zero until the activity process enters $\mathcal{B}_{k_2}$. Hence

$$\mathbb{P}\left(\tau_{k_2}(t(\nu)) > 0\right) \leq \mathbb{P}\left(T_{(k_1,l_1),(k_2,1)}(\nu) < t(\nu)\right) = \mathbb{P}\left(\frac{T_{(k_1,l_1),(k_2,1)}(\nu)}{\mathbb{E}T_{(k_1,l_1),(k_2,1)}(\nu)} < \frac{t(\nu)}{\mathbb{E}T_{(k_1,l_1),(k_2,1)}(\nu)}\right).$$

Taking the limit as $\nu \to \infty$, Theorem 3.2, gives $\lim_{\nu\to\infty} \mathbb{P}\left(\tau_{k_2}(t(\nu)) > 0\right) \leq \mathbb{P}\left(Z < \omega\right)$, and (6) follows.

7.2 Throughput near-saturation

Assume that at time $t = 0$ there is at least one user active in C_k, i.e. $X(0) = (k, l) \in \mathcal{B}_k$. Define the total full-component active time in $[0, t]$ as

$$\tau_k[0, t] := \int_0^t I_{\{X(s)=(k,L_k)\}}\, ds,$$

the residual time in C_k during $[0,t]$ as

$$R_k[0,t] := \int_0^t I_{\{X(r)\in\mathcal{B}_k\ \forall\, r\in[0,s]\}}\, ds,$$

and the full-component active time contained in the residual time in C_k during $[0,t]$ as

$$\tau_k^{\rm res}[0,t] := \int_0^t I_{\{X(r)\in\mathcal{B}_k\ \forall\, r\in[0,s]\}} I_{\{X(s)=(k,L_k)\}}\, ds.$$

For compactness, we have suppressed the implicit dependence on the parameter ν and the initial state (k,l) in the notation. From this point onwards, we will also drop the subscript k to keep the notation light.

Note that $R[0,t] \stackrel{d}{=} \min\{t, T_{(k,l),0}\}$ and that $\tau^{\rm res}[0,t] \stackrel{d}{=} \tau[0,R[0,t]]$. The random variables $\tau[0,t]$, $R[0,t]$ and $\tau^{\rm res}[0,t]$, being particular occupancy times, are non-decreasing in t on every sample path of the activity process $\{X(t)\}_{t\geq 0}$. Therefore, the random variables

$$\tau[0,\infty] := \lim_{t\to\infty} \tau[0,t], \quad R[0,\infty] := \lim_{t\to\infty} R[0,t] = T_{(k,l),0}, \quad \tau^{\rm res}[0,\infty] := \lim_{t\to\infty} \tau^{\rm res}[0,t]$$

are well defined. For $0\leq s\leq t\leq\infty$, we define

$$\tau[s,t] := \tau[0,t]-\tau[0,s], \quad R[s,t] := R[0,t]-R[0,s], \quad \tau^{\rm res}[s,t] := \tau^{\rm res}[0,t]-\tau^{\rm res}[0,s].$$

From the above definition, it is easily seen that for every sample path, $\tau^{\rm res}[s,t]$ provides a lower bound for both $\tau[s,t]$ and $R[s,t]$, as stated in the next lemma.

Lemma 7.1. *For $0\leq s\leq t\leq\infty$, $\tau^{\rm res}[s,t]\leq\tau[s,t]$ and $\tau^{\rm res}[s,t]\leq R[s,t]$.*

Proof. Rearranging terms, the differences $\tau[s,t]-\tau^{\rm res}[s,t]$ and $R[s,t]-\tau^{\rm res}[s,t]$ can both be written as integrals with an integrand that is always non-negative. □

In particular, Lemma 7.1 implies that, for every $0\leq s\leq t\leq\infty$

$$\mathbb{E}\tau^{\rm res}[s,t]\leq\mathbb{E}R[s,t]. \tag{26}$$

However, as stated in the next lemma, in the limit as $\nu\to\infty$, the ratio of the expected values of $\tau^{\rm res}[0,\infty]$ and $T_{(k,l),0}=R[0,\infty]$ converges to 1.

Lemma 7.2. *For any initial state $X(0)=(k,l)\in\mathcal{B}_k$,*

$$\lim_{\nu\to\infty}\frac{\mathbb{E}\tau^{\rm res}[0,\infty]}{\mathbb{E}T_{(k,l),0}}=1.$$

Proof. Since the ratio is clearly less than 1 by Equation (26), it suffices to show that the liminf as $\nu\to\infty$ is larger than 1. Applying the result in [23] and using (2), one obtains that for every $1\leq l\leq L_k$, if $X(0)=(k,l)$, then

$$\mathbb{E}\tau^{\rm res}[0,\infty]=\mathbb{E}\Big(\int_0^{T_{(k,l),0}} I_{\{X(s)=(k,L_k)\}}\, ds\Big)\geq\frac{1}{L_k}f_k(\nu)^{L_k-1},$$

and thus, involving (7),

$$\liminf_{\nu\to\infty}\frac{\mathbb{E}\tau^{\rm res}[0,\infty]}{\mathbb{E}T_{(k,l),0}}\geq\liminf_{\nu\to\infty}\frac{f_k(\nu)^{L_k-1}/L_k}{\mathbb{E}T_{(k,l),0}}=1.$$

□

The next proposition establishes a near-saturation property in the sense that if $X(0)=(k,l)\in B_k$, then for any time period $t(\nu)=o(\mathbb{E}T_{(k,l),0})$ every user in C_k will be active an arbitrarily large fraction of the time with probability one as $\nu\to\infty$.

Proposition 7.3. *Suppose that $X(0)=(k,l)\in\mathcal{B}_k$ and that $T_{(k,l),0}/\mathbb{E}T_{(k,l),0}\stackrel{d}{\to}Z$ as $\nu\to\infty$. Then for every $\omega\in[0,1]$ and every $\delta>0$,*

$$\liminf_{\nu\to\infty}\mathbb{P}\left(\tau[0,\omega\mathbb{E}T_{(k,l),0}]\geq(1-\delta)\omega\mathbb{E}T_{(k,l),0}\right)\geq\mathbb{P}\left(Z\geq\omega\right).$$

In particular, for any $t(\nu)=o(\mathbb{E}T_{(k,l),0}(\nu))$, $\liminf_{\nu\to\infty}\mathbb{P}\left(\tau[0,t(\nu)]\geq(1-\delta)t(\nu)\right)=1$.

Remark 1. As mentioned earlier, the hypothesis that $T_{(k,l),0}/\mathbb{E}T_{(k,l),0} \xrightarrow{d} Z$ is not just a convenient assumption, but something that we actually know. In particular, Proposition 4.6 says that $Z \stackrel{d}{=} \text{Exp}(1)$. Moreover, since the result holds for every initial state in $\mathcal{B}_k$, it is true also for a random initial state in $\mathcal{B}_k$. Indeed, as seen in Section 4, the convergence in distribution of $T_{(k,l),0}(\nu)/\mathbb{E}T_{(k,l),0}(\nu)$ to Z as $\nu \to \infty$ does not depend on the initial state, as long as it belongs to $\mathcal{B}_k$.

Proof. In order to keep the notation light, we denote in this proof the hitting time $T_{(k,l),0}$ by T. Firstly, Lemma 7.1 implies that

$$\mathbb{P}\left(\tau[0,\omega\mathbb{E}T] \geq (1-\delta)\omega\mathbb{E}T\right) \geq \mathbb{P}\left(\tau^{\text{res}}[0,\omega\mathbb{E}T] \geq (1-\delta)\omega\mathbb{E}T_0\right).$$

Moreover, by definition of $R[0,t] = \min\{t, T\}$, we have

$$\mathbb{P}\left(R[0,\omega\mathbb{E}T] \geq \omega\mathbb{E}T\right) = \mathbb{P}\left(T \geq \omega\mathbb{E}T\right).$$

In view of the hypothesis that $T(\nu)/\mathbb{E}T(\nu) \xrightarrow{d} Z$ as $\nu \to \infty$, it therefore suffices to prove that for every $\omega \in [0,1]$ and every $\delta > 0$,

$$\liminf_{\nu\to\infty} \mathbb{P}\left(\tau^{\text{res}}[0,\omega\mathbb{E}T] \geq (1-\delta)\omega\mathbb{E}T\right) \geq \liminf_{\nu\to\infty} \mathbb{P}\left(R[0,\omega\mathbb{E}T] \geq \omega\mathbb{E}T\right).$$

Suppose that this latter statement is false, i.e. there exist $\omega_0 \in [0,1]$, $\delta > 0$ and $\eta > 0$ such that

$$\liminf_{\nu\to\infty} \mathbb{P}\left(\tau^{\text{res}}[0,\omega_0\mathbb{E}T] \geq (1-\delta)\omega_0\mathbb{E}T\right) \leq \liminf_{\nu\to\infty} \mathbb{P}\left(R[0,\omega_0\mathbb{E}T] \geq \omega_0\mathbb{E}T\right) - \eta. \tag{27}$$

Then it can be shown that there exists $\varepsilon_{\omega_0,\delta} > 0$ such that

$$\liminf_{\nu\to\infty} \frac{\mathbb{E}\tau^{\text{res}}[0,\omega_0\mathbb{E}T]}{\mathbb{E}T} \leq \liminf_{\nu\to\infty} \frac{\mathbb{E}R[0,\omega_0\mathbb{E}T]}{\mathbb{E}T} - \varepsilon_{\omega_0,\delta}. \tag{28}$$

Indeed,

$$\begin{aligned}
\liminf_{\nu\to\infty} \left(\frac{\mathbb{E}R[0,\omega_0\mathbb{E}T]}{\mathbb{E}T} - \frac{\mathbb{E}\tau^{\text{res}}[0,\omega_0\mathbb{E}T]}{\mathbb{E}T}\right) &= \liminf_{\nu\to\infty} \int_0^\infty \mathbb{P}\left(\frac{R[0,\omega_0\mathbb{E}T]}{\mathbb{E}T} \geq y\right) - \mathbb{P}\left(\frac{\tau^{\text{res}}[0,\omega_0\mathbb{E}T]}{\mathbb{E}T} \geq y\right) dy \\
&\overset{\text{Lemma 7.1}}{\geq} \liminf_{\nu\to\infty} \int_{(1-\delta)\omega_0}^{\omega_0} \mathbb{P}\left(\frac{R[0,\omega_0\mathbb{E}T]}{\mathbb{E}T} \geq y\right) - \mathbb{P}\left(\frac{\tau^{\text{res}}[0,\omega_0\mathbb{E}T]}{\mathbb{E}T} \geq y\right) dy \\
&\geq \int_{(1-\delta)\omega_0}^{\omega_0} \liminf_{\nu\to\infty} \left(\mathbb{P}\left(\frac{R[0,\omega_0\mathbb{E}T]}{\mathbb{E}T} \geq y\right) - \mathbb{P}\left(\frac{\tau^{\text{res}}[0,\omega_0\mathbb{E}T]}{\mathbb{E}T} \geq y\right)\right) dy \\
&\geq \eta\delta\omega_0 > 0,
\end{aligned}$$

where the second last inequality follows from the generalized Fatou's lemma, while the last inequality follows from (27) and is illustrated by Figure 5. Thus we can take $\varepsilon_{\omega_0,\delta} := \eta\delta\omega_0$.

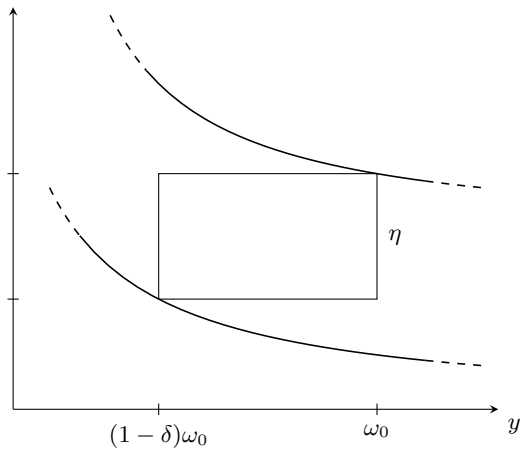


Figure 5: $\mathbb{P}\left(\tau^{\text{res}}[0,\omega_0\mathbb{E}T]/\mathbb{E}T \geq y\right)$ (lower line) vs $\mathbb{P}\left(R[0,\omega_0\mathbb{E}T]/\mathbb{E}T \geq y\right)$ (upper line)

Equation (26) yields

$$\liminf_{\nu\to\infty} \frac{\mathbb{E}\tau^{\text{res}}[\omega_0\mathbb{E}T,\infty]}{\mathbb{E}T} \leq \liminf_{\nu\to\infty} \frac{\mathbb{E}R[\omega_0\mathbb{E}T,\infty]}{\mathbb{E}T}, \tag{29}$$

and thus, summing term by term (28) and (29), since by definition $\mathbb{E}R[0,\infty] = \mathbb{E}T$,

$$\liminf_{\nu\to\infty} \frac{\mathbb{E}\tau^{\mathrm{res}}[0,\infty]}{\mathbb{E}T} \leq \liminf_{\nu\to\infty} \frac{\mathbb{E}R[0,\infty]}{\mathbb{E}T} - \varepsilon_{\omega_0,\delta} = 1 - \varepsilon_{\omega_0,\delta},$$

which contradicts Lemma 7.2. □

8 Mixing times

In the previous sections we have analyzed the transient behavior of the Markov process $\{X(t)\}_{t\geq 0}$ in terms of hitting times and we have shown how this leads to starvation of individual users over certain time scales. In this section we turn attention to the long-run behavior of the Markov process $\{X(t)\}_{t\geq 0}$ and in particular examine the rate of convergence to the stationary distribution. We measure the rate of convergence in terms of the total variation distance and the so-called *mixing time*, which describes the time required for the distance to stationarity to become small.

The mixing time becomes particularly relevant when the network has two or more dominant components which together attract the entire probability mass in the limit as $\nu \to \infty$. Indeed, in this case, the mixing time provides an indication of how long it takes the activity process to reach a certain level of fairness among the dominant components. We will prove a lower bound for the mixing time using the notion of conductance.

The maximal distance over $x \in \Omega$, measured in terms of total variation, between the distribution at time t and the stationary distribution is defined as

$$d(t,\nu) := \max_{x\in\Omega} \|\mathbb{P}\left(X(t) \in \cdot \,|X(0) = x\right) - \pi(\nu)\|_{\mathrm{TV}}.$$

We define the mixing time of the process $\{X(t)\}_{t\geq 0}$ as

$$t_{\mathrm{mix}}(\varepsilon,\nu) := \inf\{t \geq 0 : d(t,\nu) \leq \varepsilon\}.$$

For a fixed $r \in (0,1)$ consider the subset $\tilde{K}(r)$ of branches whose stationary probability is asymptotically no more than r, i.e. $\tilde{K}(r) := \{k \,:\, \lim_{\nu\to\infty} \pi_{\mathcal{B}_k}(\nu) \leq r\}$. Define $\kappa \in \tilde{K}(r)$ as the index corresponding to the branch $\mathcal{B}_\kappa$ which has asymptotically the largest mean escape time, i.e. such that for every $j \in \tilde{K}(r)$,

$$\lim_{\nu\to\infty} \frac{\mathbb{E}T_{(\kappa,1),0}(\nu)}{\mathbb{E}T_{(j,1),0}(\nu)} = \lim_{\nu\to\infty} \frac{L_j f_\kappa(\nu)^{L_\kappa - 1}}{L_\kappa f_j(\nu)^{L_j - 1}} \geq 1. \tag{30}$$

The next result shows that the mixing time is asymptotically at least of the same order of magnitude as the escape time from branch $\mathcal{B}_\kappa$.

Proposition 8.1. *For any $r \in (0,1)$ and $\varepsilon \in \left(0, \frac{1-r}{2}\right)$, there exists a constant $C_{\varepsilon,r} > 0$ and $\nu_0 > 0$ such that for every $\nu > \nu_0$*

$$t_{\mathrm{mix}}(\varepsilon,\nu) \geq C_{\varepsilon,r} \frac{f_\kappa(\nu)^{L_\kappa - 1}}{L_\kappa}.$$

Proposition 8.1 shows that it can take an extremely long time for the process $\{X(t)\}_{t\geq 0}$ to reach stationarity, especially when ν is large. Such a long mixing time is typically due to the activity process being stuck for a considerable period in one of the components, and thus not visiting the states in the other components. We will prove Proposition 8.1 exploiting the presence of a bottleneck in the state space and using the notion of conductance.

For $S \subseteq \Omega$, let $\pi_S(\nu) = \sum_{(k,l)\in S} \pi_{(k,l)}(\nu)$ be the stationary probability of S. Define the *probability flow out of S* as

$$Q(S,S^c) := \sum_{(k,l)\in S,(j,m)\in S^c} \pi_{(k,l)}(\nu) q((k,l),(j,m))$$

and its *conductance* as $\Phi(S) := Q(S,S^c)/\pi_S$. All the quantities we just defined clearly depend on ν, but we suppress it for compactness. The *conductance profile* of the process $\{X(t)\}_{t\geq 0}$ is defined as

$$\Phi_r(\nu) := \min_{S\,:\,\pi(S)\leq r} \Phi(S).$$

The following result, valid for any Markov process on a finite state space Ω with conductance profile Φ_r, shows how the conductance of the process yields a lower bound on the mixing time. It is a continuous-time version of Theorem 7.3 in [21] and the proof is relegated to D.

Lemma 8.2. *For any $r \in (0,1)$ and any $\varepsilon \in \left(0, \frac{1-r}{2}\right)$,*

$$t_{\text{mix}}(\varepsilon, \nu) \geq \frac{1-r-2\varepsilon}{\Phi_r(\nu)}.$$

In order to get a sharp bound for the conductance and hence a sharp lower bound for the mixing time, we need to identify a subset S with low conductance. As proved in [19], it suffices to look at the connected subsets of the state space. Therefore the branches in $\tilde{K}(r)$ become naturally good candidates for being the lowest-conductance subsets of Ω. From (2) it follows that if (k,l) and (k,m) are states in the same branch $\mathcal{B}_k$, then

$$\frac{\pi_{(k,m)}(\nu)}{\pi_{(k,l)}(\nu)} = \frac{l!(L_k-l)!}{m!(L_k-m)!} f(\nu)^{m-l}, \quad \text{as } \nu \to \infty.$$

Thus the conductance of $\mathcal{B}_k$ satisfies

$$\Phi(\mathcal{B}_k) = \frac{\pi_{(k,1)}(\nu)\cdot 1}{\sum_{l=1}^{L_k} \pi_{(k,l)}(\nu)} = \frac{\frac{\pi_{(k,1)}(\nu)}{\pi_{(k,L_k)}(\nu)}}{\sum_{l=1}^{L_k} \frac{\pi_{(k,l)}(\nu)}{\pi_{(k,L_k)}(\nu)}} \sim L_k f_k(\nu)^{1-L_k}, \quad \text{as } \nu \to \infty.$$

Thanks to the definition of κ, $\mathcal{B}_\kappa$ is asymptotically the branch with the smallest conductance, see Equation (30). Since by definition $\Phi_r(\nu) \leq \Phi(\mathcal{B}_\kappa)$, Lemma 8.2 then gives that for every $\varepsilon \in (0, \frac{1}{4})$ and for ν sufficiently large

$$t_{\text{mix}}(\varepsilon, \nu) \geq (1-r-2\varepsilon)\frac{f_\kappa(\nu)^{L_\kappa - 1}}{L_\kappa},$$

which completes the proof of the lower bound claimed in Proposition 8.1.

9 Conclusions

We have studied hitting times and mixing properties in dense wireless random-access networks. We have represented the activity processes in such networks in terms of Markov processes on complete partite graphs. In particular, in dense networks, high activity rates lead to network behavior in which users in maximum independent sets coalesce into components which compete for the wireless medium. We have shown that components monopolize the wireless medium for extremely long periods, which leads to long mixing times and starvation of all other components. Hence, users in a particular component alternate between enjoying long periods of access, and facing long periods of starvation.

While the slow nature of the transitions is a common characteristic, the asymptotic distribution of the scaled transition time depends crucially on the structure of the network and on the initial and target components and there is a notable variety of possible scenarios. In particular, in some scenarios, the distribution of the scaled transition time is non-exponential. This is due to the heterogeneous activation rates among components in conjunction with the presence of intermediate components where the activity process resides for long periods along the transition.

The complete partite graphs that we focused on in the present paper, are arguably the worst possible networks in terms of transition times and starvation effects, given the size of the network. Indeed, the fact that the users are grouped into components, with no interference within components and full interference between components, turns out to be a key element for starvation to occur. This is reflected in the fact that the transition times exhibit exponential growth in the component size. Graphs that are non-partite, or partite but not complete, will have a less extreme tendency for starvation, although the issue may still arise to a milder degree. For example, interference between nodes inside the same component will reduce the size of the maximum independent set, and lower the likelihood of the maximal activity states relative to the bottleneck state where all the nodes are inactive. This is borne out by the model extensions in Section 6 where the order-of-magnitude of the transition time is governed by the maximum independent set within a component. Likewise, lack of interference between nodes in different components will result in bottleneck states where some of the nodes may be active, and raise the likelihood of the bottleneck state relative to the dominant activity states. This is illustrated by the work in [29] which investigates the transition times and delays in a toric grid. The toric grid is a bi-partite graph, but with fewer edges between the two components. The results in [29] show that the delays and long transition times, while still severe, are of a lower order than for the complete bi-partite graph.

References

[1] D. J. Aldous. Markov chains with almost exponential hitting times. *Stochastic Processes and Their Applications*, 13(3):305--310, 1982.

[2] D. J. Aldous and M. Brown. Inequalities for rare events in time-reversible Markov chains I. In *Stochastic Inequalities*, volume 22 of *Lecture Notes Monograph Series*, pages 1--16. IMS, 1992.

[3] D. J. Aldous and M. Brown. Inequalities for rare events in time-reversible Markov chains II. *Stochastic Processes and Their Applications*, 44(1):15--25, 1993.

[4] R. R. Boorstyn, A. Kershenbaum, B. Maglaris, and V. Sahin. Throughput analysis in multihop CSMA packet radio networks. *IEEE Transactions on Communications*, 35:267--274, 1987.

[5] R. Gouet and A. Sudbury. Blocking and dimer processes on the Cayley tree. *Journal of Statistical Physics*, 130(5):935--955, 2008.

[6] R. Grübel and M. Reich. Rarity and exponentiality: an extension of Keilson's theorem, with applications. *Journal of Applied Probability*, 42(2):393--406, 2005.

[7] O. Häggström. Ergodicity of the hard-core model on $\mathbb{Z}^2$ with parity-dependent activities. *Arkiv för Matematik*, 35(1):171--184, 1997.

[8] L. Jiang and J. Walrand. A distributed CSMA algorithm for throughput and utility maximization in wireless networks. *IEEE Transaction on Networking*, 18(3):960--972, 2010.

[9] O. Jouini and Y. Dallery. Moments of first passage times in general birthdeath processes. *Mathematical Methods of Operations Research*, 68(1):49--76, 2008.

[10] S. Karlin and J. McGregor. Coincidence properties of birth and death processes. *Pacific Journal of Mathematics*, 9:1109--1140, 1959.

[11] J. Keilson. A review of transient behavior in regular diffusion and birth-death processes. Part II. *Journal of Applied Probability*, 2(2):405--428, 1965.

[12] J. Keilson. A limit theorem for passage times in ergodic regenerative processes. *The Annals of Mathematical Statistics*, 37(4):866--870, 1966.

[13] J. Keilson. Log-concavity and log-convexity in passage time densities of diffusion and birth-death processes. *Journal of Applied Probability*, 8(2):391--398, 1971.

[14] J. Keilson. *Markov Chain Models--Rarity and Exponentiality*. Springer, 1979.

[15] F. P. Kelly. *Reversibility and Stochastic Networks*. Wiley, 1979.

[16] F. P. Kelly. Stochastic models of computer communication systems. *Journal of the Royal Statistical Society. Series B. Methodological*, 47 (3):379--395, 1985.

[17] A. Kershenbaum, R. R. Boorstyn, and M. S. Chen. An algorithm for evaluation of throughput in multihop packet radio networks with complex topologies. *IEEE Journal on Selected Areas in Communications*, 5(6):1003--1012, 1987.

[18] M. Koseoglu, E. Karasan, and M. Alanyali. Effect of network density and size on the short-term fairness performance of CSMA systems. *EURASIP Journal on Wireless Communications and Networking*, 2012(1):1--15, 2012.

[19] G. F. Lawler and A. D. Sokal. Bounds on the L^2 spectrum for Markov chains and Markov processes: a generalization of Cheeger's inequality. *Transactions of the American Mathematical Society*, 309(2):557--580, 1988.

[20] W. Ledermann and G. E. H. Reuter. Spectral theory for the differential equations of simple birth and death processes. *Philosophical Transactions of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 246:321--369, 1954.

[21] D. Levin, Y. Peres, and E. Wilmer. *Markov Chains and Mixing Times*. AMS, 2009.

[22] E. Olivieri and M. E. Vares. *Large Deviations and Metastability*. Cambridge University Press, 2005.

[23] V. T. Stefanov and S. Wang. A note on integrals for birth--death processes. *Mathematical Biosciences*, 168(2):161--165, 2000.

[24] Y. Suhov and U. A. Rozikov. A hard-core model on a Cayley tree: an example of a loss network. *Queueing Systems*, 46(1):197--212, 2004.

[25] P. M. van de Ven, A. J. E. M. Janssen, J. S. H. van Leeuwaarden, and S. C. Borst. Achieving target throughputs in random-access networks. *Performance Evaluation*, 68(11):1103--1117, 2011.

[26] X. Wang and K. Kar. Throughput modeling and fairness issues in CSMA/CA based ad-hoc networks. In *Proc. IEEE INFOCOM Conf.*, volume 1, pages 23--34, 2005.

[27] S. Zachary and I. Ziedins. Loss networks and Markov random fields. *Journal of Applied Probability*, 36(2):403--414, 1999.

[28] A. Zocca, S. C. Borst, and J. S. H. van Leeuwaarden. Mixing properties of CSMA networks on partite graphs. In *Proc. ValueTools Conf.*, pages 117--126, 2012.

[29] A. Zocca, S. C. Borst, J. S. H. van Leeuwaarden, and F. R. Nardi. Delay performance in random-access grid networks. *Performance Evaluation*, 70(10):900--915, 2013.

A Proof of Lemma 4.3

We will prove a slightly more general version of Lemma 4.3, assuming the rates of the process are those described in Subsection 4.3. Order the state space as $\Omega = \{L, L-1, \ldots, 1, 0\}$ and consider the generator matrix $Q(\nu)$ of the process $\{X(t)\}_{t\geq 0}$ with 0 an absorbing state. That is,

$$Q_L(\nu) = \begin{bmatrix} q_L & d_L & 0 & & \cdots & & 0 \\ a_{L-1}f(\nu) & q_{L-1} & d_{L-1} & 0 & & & \\ 0 & a_{L-2}f(\nu) & q_{L-2} & d_{L-2} & 0 & & \\ \vdots & & & \ddots & \ddots & \ddots & & \\ & & & & 0 & a_1 f(\nu) & q_1 & d_1 \\ 0 & & \cdots & & 0 & 0 & 0 & 0 \end{bmatrix},$$

where the diagonal elements are $q_l(\nu) = -a_l f(\nu) - d_l$ for $l = 1, \ldots, L_1$ and $q_L(\nu) = -d_L$. The matrix $Q(\nu)$ can be written as

$$Q(\nu) = \begin{pmatrix} \mathbf{T}(\nu) & \mathbf{t}(\nu) \\ \mathbf{0} & 0 \end{pmatrix},$$

where $\mathbf{T}(\nu)$ is an $L \times L$ invertible matrix. Since the characteristic polynomials of $-Q(\nu)$ and $-\mathbf{T}(\nu)$ satisfy the relation $p_{-Q(\nu)}(z) = -z\, p_{-\mathbf{T}(\nu)}(z)$, the spectrum of $-Q(\nu)$ consists of that of $-\mathbf{T}(\nu)$ plus the eigenvalue zero with multiplicity one.
Denote by $D(\nu)$ the $L \times L$ diagonal matrix, whose diagonal entries are $\{\sqrt{\xi_l(\nu)}\}_{i=L}^{1}$, where the ξ's are the so-called *potential coefficients*, defined recursively as

$$\xi_L(\nu) = 1 \quad \text{and} \quad \xi_{l-1}(\nu) = \frac{d_l}{a_l f(\nu)}\xi_l(\nu), \quad l = L-1, \ldots, 1.$$

The $L \times L$ matrix $G(\nu) = -D(\nu)^{1/2}\mathbf{T}(\nu)D(\nu)^{-1/2}$ is tridiagonal and symmetric with diagonal entries $G_{l,l}(\nu) = q_{L-l+1}(\nu)$ and $G_{l,l+1}(\nu) = g_{l+1,l}(\nu) = -\sqrt{d_l a_{l-1} f(\nu)}$. Since $G(\nu)$ is similar to $-\mathbf{T}(\nu)$, they have the same spectrum. Denote by $\mathcal{D}(p, R)$ the closed disc centered in p with radius R, i.e. $\mathcal{D}(p, R) = \{z \in \mathbb{C} \ : \ |z - p| \leq R\}$. Consider the

Gershgorin discs $\{\mathcal{D}_l(\nu)\}_{l=1}^{L}$ of $G(\nu)$, defined as $\mathcal{D}_l(\nu) := \mathcal{D}(-q_l(\nu), R_l(\nu))$, where the radius $R_l(\nu)$ is the sum of the absolute values of the non-diagonal entries in the $L-l+1$-th row, i.e. $R_l(\nu) := \sum_{m\neq L-l+1} |G_{L-l+1,m}(\nu)|$. Then

$$\begin{aligned}
\mathcal{D}_L(\nu) &= \mathcal{D}(d_L, \sqrt{d_L a_{L-1} f(\nu)}),\\
\mathcal{D}_{L-1}(\nu) &= \mathcal{D}(d_{L-1} + a_{L-1} f(\nu), \sqrt{d_L a_{L-1} f(\nu)} + \sqrt{d_{L-1} a_{L-2} f(\nu)}),\\
&\vdots\\
\mathcal{D}_2(\nu) &= \mathcal{D}(d_2 + a_2 f(\nu), \sqrt{d_3 a_2 f(\nu)} + \sqrt{d_2 a_1 f(\nu)}),\\
\mathcal{D}_1(\nu) &= \mathcal{D}(d_1 + a_1 f(\nu), \sqrt{d_2 a_1 f(\nu)}).
\end{aligned}$$

We now exploit the second Gershgorin circle theorem, which is reproduced here for completeness.

Theorem. *If the union of j Gershgorin discs of a real $r \times r$ matrix A is disjoint from the union of the other $r-j$ Gershgorin discs, then the former union contains exactly j and the latter the remaining $r-j$ eigenvalues of A.*

In our case, for ν sufficiently large, the disc $\mathcal{D}_L(\nu)$ does not intersect with the union $\bigcup_{l=1}^{L-1} \mathcal{D}_l(\nu)$, thus the smallest eigenvalue $\theta_1(\nu)$ lies in $\mathcal{D}_L(\nu)$ and the other $L-1$ ones in $\bigcup_{l=1}^{L-1} \mathcal{D}_l(\nu)$. Hence, for ν sufficiently large, the following inequalities hold

$$\begin{aligned}
&\theta_1(\nu) \leq A + B\sqrt{f(\nu)},\\
&\theta_i(\nu) \geq Cf(\nu) - D\sqrt{f(\nu)}, \quad i = 2, \ldots, L,
\end{aligned}$$

where $A, B, C, D \in \mathbb{R}_+$ and, more precisely, $A = d_L$, $B = \sqrt{d_L a_{L-1}}$, $C = \min_{l=1,\ldots,L-1} a_l$ and $D = \max\{\sqrt{d_L a_{L-1}} + \sqrt{d_{L-1} a_{L-2}}, \ldots, \sqrt{d_3 a_2} + \sqrt{d_2 a_1}, \sqrt{d_2 a_1}\}$. Therefore, for ν sufficiently large,

$$0 < \frac{\theta_1(\nu)}{\theta_i(\nu)} \leq \frac{A + B\sqrt{f(\nu)}}{Cf(\nu) - D\sqrt{f(\nu)}},$$

and so $\lim_{\nu\to\infty} \theta_1(\nu)/\theta_i(\nu) = 0$ for $i = 2, \ldots, L$. Hence,

$$\mathbb{E}T_{L,0}(\nu) \cdot \theta_1(\nu) = 1 + \sum_{i=2}^{L} \frac{\theta_1(\nu)}{\theta_i(\nu)} \to 1, \quad \text{as } \nu \to \infty,$$

while for $2 \leq i \leq L$,

$$\mathbb{E}T_{L,0}(\nu) \cdot \theta_i(\nu) > \frac{\theta_i(\nu)}{\theta_1(\nu)} \to \infty, \quad \text{as } \nu \to \infty.$$

B Proof of Lemma 4.5

(a) The proof consists of a lower and an upper bound which asymptotically coincide. Indeed, using the bounds in Property (ii), one obtains that

$$\liminf_{\nu\to\infty} \frac{\mathbb{E}T(\nu)}{\mathbb{E}U(\nu)} \geq \liminf_{\nu\to\infty} \frac{\mathbb{E}U(\nu) - \mathbb{E}V(\nu)}{\mathbb{E}U(\nu)} = 1 \quad \text{and} \quad \limsup_{\nu\to\infty} \frac{\mathbb{E}T(\nu)}{\mathbb{E}U(\nu)} \leq \limsup_{\nu\to\infty} \frac{\mathbb{E}U(\nu) + \mathbb{E}W(\nu)}{\mathbb{E}U(\nu)} = 1.$$

(b) Once again, the proof consists of a lower and an upper bound which asymptotically coincide for all the continuity points of the tail distribution of Z, which will be denoted by $F(s) = \mathbb{P}(Z > s)$.

For the lower bound, argue as follows. Property (i) implies that for any $\delta \in (0,1)$, $\mathbb{E}W(\nu) \leq \delta \mathbb{E}U(\nu)$ for ν sufficiently large. Thus, using the lower bound in Property (iii), for ν sufficiently large,

$$\begin{aligned}
\mathbb{P}\left(\frac{T(\nu)}{\mathbb{E}T(\nu)} > t\right) &\geq \mathbb{P}\left(\frac{U(\nu) - V(\nu)}{\mathbb{E}U(\nu) + \mathbb{E}W(\nu)} > t\right)\\
&\geq \mathbb{P}\left(U(\nu) - V(\nu) > \mathbb{E}U(\nu)(1+\delta)t\right)\\
&\geq \mathbb{P}\left(U(\nu) > \mathbb{E}U(\nu)(1+2\delta)t\right) - \mathbb{P}\left(V(\nu) > \delta \mathbb{E}U(\nu)t\right).
\end{aligned}$$

Property (iii) implies that

$$\lim_{\nu\to\infty} \mathbb{P}\left(\frac{U(\nu)}{\mathbb{E}U(\nu)} > (1+2\delta)t\right) = F\big((1+2\delta)\,t\big).$$

Property (i) implies that for ν sufficiently large, $\mathbb{E}V(\nu) \leq \delta^2 \mathbb{E}U(\nu)$, so that

$$\mathbb{P}\left(V(\nu) > \delta \mathbb{E}U(\nu)t\right) \leq \frac{\mathbb{E}V(\nu)}{\delta \mathbb{E}U(\nu)t} \leq \frac{\delta}{t},$$

by Markov's inequality. Taking liminf's, we obtain

$$\liminf_{\nu\to\infty} \mathbb{P}\left(\frac{T(\nu)}{\mathbb{E}T(\nu)} > t\right) \geq F\big((1+2\delta)t\big) - \frac{\delta}{t}.$$

Letting $\delta \downarrow 0$, we find

$$\liminf_{\nu\to\infty} \mathbb{P}\left(\frac{T(\nu)}{\mathbb{E}T(\nu)} > t\right) \geq F\left(t\right). \tag{31}$$

For the upper bound, argue as follows. Property (i) implies that for any $\delta \in (0,1)$, $\mathbb{E}V(\nu) \leq \delta \mathbb{E}U(\nu)$ for ν sufficiently large. Thus, using the upper bound in Property (iii), for ν sufficiently large,

$$\begin{aligned}
\mathbb{P}\left(\frac{T(\nu)}{\mathbb{E}T(\nu)} > t\right) &\leq \mathbb{P}\left(\frac{U(\nu)+W(\nu)}{\mathbb{E}U(\nu)-\mathbb{E}V(\nu)} > t\right) \\
&\leq \mathbb{P}\left(U(\nu)+W(\nu) > (1-\delta)\mathbb{E}U(\nu)t\right) \\
&\leq \mathbb{P}\left(U(\nu) > (1-2\delta)\mathbb{E}U(\nu)t\right) + \mathbb{P}\left(W(\nu) > \delta \mathbb{E}U(\nu)t\right).
\end{aligned}$$

Property (iii) implies that

$$\lim_{\nu\to\infty} \mathbb{P}\left(\frac{U(\nu)}{\mathbb{E}U(\nu)} > (1-2\delta)t\right) = F\big((1-2\delta)t\big).$$

Property (i) implies that for ν sufficiently large, $\mathbb{E}W(\nu) \leq \delta^2 \mathbb{E}U(\nu)$, so that

$$\mathbb{P}\left(W(\nu) > \delta \mathbb{E}U(\nu)t\right) \leq \frac{\mathbb{E}W(\nu)}{\delta \mathbb{E}U(\nu)t} \leq \frac{\delta}{t},$$

by Markov's inequality. Taking limsup's, we obtain

$$\limsup_{\nu\to\infty} \mathbb{P}\left(\frac{T(\nu)}{\mathbb{E}T(\nu)} > t\right) \leq F\big((1-2\delta)\,t\big) + \frac{\delta}{t}.$$

Letting $\delta \downarrow 0$, we find

$$\limsup_{\nu\to\infty} \mathbb{P}\left(\frac{T(\nu)}{\mathbb{E}T(\nu)} > t\right) \leq F\left(t\right). \tag{32}$$

Combining (31) and (32) completes the proof.

C Proof of Lemma 5.3

Statement (a) is trivial, since $0 \leq \mathcal{L}_{T_k(\nu)/\mathbb{E}T_k(\nu)}(s) \leq 1$ for all $s \in [0,\infty[$ and $\beta_k = \lim_{\nu\to\infty} \mathbb{E}N_k(\nu) = 0$ for every $k \in \mathcal{N}$. Statement (b) follows immediately after substituting $\beta_k \in \mathbb{R}_+$ and $\gamma_k \in (0,1]$ in the limit, since $T_k(\nu)/\mathbb{E}T_k(\nu) \xrightarrow{d} \text{Exp}(1)$. The proof of claim (c) is more involved and we need an auxiliary lemma.

Let $S_k(\nu) := \sum_{i=1}^{N_k(\nu)} T_k^{(i)}(\nu)$. From the integrability of $N_k(\nu)$ and $T_k(\nu)$ it follows that $\mathbb{E}S_k(\nu) = \mathbb{E}N_k(\nu) \cdot \mathbb{E}T_k(\nu)$. Consider the random variable

$$\tilde{S}(\nu) = g(\nu) \sum_{i=1}^{N_k(\nu)} \frac{T_k^{(i)}(\nu)}{\mathbb{E}T_k^{(i)}(\nu)} = g(\nu)\mathbb{E}N_k(\nu)\frac{S_k(\nu)}{\mathbb{E}S_k(\nu)}.$$

Since $N_k(\nu)$ has a geometric distribution, the Laplace transform of $S_k(\nu)/\mathbb{E}S_k(\nu)$ is given by

$$\mathcal{L}_{S_k(\nu)/\mathbb{E}S_k(\nu)}(s) = G_{N_k(\nu)}\left(\mathcal{L}_{T_k(\nu)/\mathbb{E}T_k(\nu)}(s/\mathbb{E}N_k(\nu))\right) = \frac{1}{1+(1-\mathcal{L}_{T_k(\nu)/\mathbb{E}T_k(\nu)}(s/\mathbb{E}N_k(\nu)))\cdot\mathbb{E}N_k(\nu)},$$

and hence

$$\mathcal{L}_{\tilde{S}(\nu)}(s) = \mathcal{L}_{S_k(\nu)/\mathbb{E}S_k(\nu)}(sg(\nu)\mathbb{E}N_k(\nu)) = \frac{1}{1+(1-\mathcal{L}_{T_k(\nu)/\mathbb{E}T_k(\nu)}(sg(\nu)))\cdot\mathbb{E}N_k(\nu)}. \tag{33}$$

One can check that, if $k \in \mathcal{S}$, then

(1) $N_k(\nu)/\mathbb{E}N_k(\nu) \xrightarrow{d} \text{Exp}(1)$, and

(2) $\lim_{\nu\to\infty} \frac{\text{Var}(T_k(\nu)/\mathbb{E}T_k(\nu))}{\mathbb{E}N_k(\nu)} = 0$.

Fact (1) is a standard result for geometric random variables which uses only the fact that $\lim_{\nu\to\infty}\mathbb{E}N_k(\nu) = \infty$ for $k \in \mathcal{S}$. Moreover, since $T_k(\nu)$ is a first passage time of a birth-and-death process, using Corollary 4 in [9], we can obtain explicitly the asymptotic order-of-magnitude of $\text{Var}(T_k(\nu))$ and prove that $\lim_{\nu\to\infty}\text{Var}(T_k(\nu))/(\mathbb{E}T_k(\nu))^2 = 1$, from which claim (2) follows.

These two facts and the technical lemma below imply that $S_k(\nu)/\mathbb{E}S_k(\nu) \xrightarrow{d} \text{Exp}(1)$, and hence

$$\lim_{\nu\to\infty}\mathcal{L}_{\tilde{S}(\nu)}(s) = \lim_{\nu\to\infty}\mathcal{L}_{S_k(\nu)/\mathbb{E}S_k(\nu)}(sg(\nu)\mathbb{E}N_k(\nu)) = \frac{1}{1+s\gamma_k}.$$

This fact, together with Equation (33), implies that

$$\lim_{\nu\to\infty}(1-\mathcal{L}_{T_k(\nu)/\mathbb{E}T_k(\nu)}(sg(\nu)))\cdot\mathbb{E}N_k(\nu) = s\gamma_k,$$

and the proof of statement (c) is concluded. We now state and prove the technical lemma mentioned above.

Lemma. *Assume that*

(i) $\{X_i(\nu)\}_{i\geq 0}$ *is a sequence of i.i.d. integrable random variables,* $X_i(\nu) \stackrel{d}{=} X(\nu)$, *for every* $\nu > 0$,

(ii) $N(\nu)$ *is an integer-valued random variable, independent of all the* $X_i(\nu)$*'s and integrable for every* $\nu > 0$,

(iii) $N(\nu)/\mathbb{E}N(\nu) \xrightarrow{d} Z$ *as* $\nu\to\infty$, *with* $\mathbb{P}(Z=0) = 0$,

(iv) $\lim_{\nu\to\infty}\frac{\text{Var}(X(\nu)/\mathbb{E}X(\nu))}{\mathbb{E}N(\nu)} = 0$.

Define $S(\nu) := \sum_{i=1}^{N(\nu)} X_i(\nu)$. *Then*

$$\frac{S(\nu)}{\mathbb{E}S(\nu)} \xrightarrow{d} Z, \qquad \text{as } \nu\to\infty.$$

Proof. Firstly, Wald's identity guarantees that $S_N(\nu)$ is integrable and that $\mathbb{E}S(\nu) = \mathbb{E}N(\nu)\mathbb{E}X(\nu)$ for every $\nu > 0$. Hence, without loss of generality, we can assume that $\mathbb{E}X(\nu) = 1$ and study the asymptotic distribution of $S(\nu)/\mathbb{E}N(\nu)$. Define $S_n(\nu) := \sum_{i=1}^n X_i(\nu)$. Note that we can rewrite

$$\frac{S(\nu)}{\mathbb{E}N(\nu)} = \frac{S(\nu)}{N(\nu)}\frac{N(\nu)}{\mathbb{E}N(\nu)}.$$

Now we claim that

$$\frac{S(\nu)}{N(\nu)} \xrightarrow{\mathbb{P}} 1, \qquad \text{as } \nu\to\infty. \tag{34}$$

Indeed, for every $\epsilon > 0$ we may write

$$\begin{aligned}
\mathbb{P}(|S(\nu)/N(\nu) - \mathbb{E}X(\nu)| > \delta) &= \sum_{n=0}^{\infty}\mathbb{P}(|S(\nu)/N(\nu) - \mathbb{E}X(\nu)| > \delta,\, N(\nu) = n) \\
&= \sum_{n=0}^{\lfloor\epsilon\mathbb{E}N(\nu)\rfloor}\mathbb{P}(\dots) + \sum_{n=\lfloor\epsilon\mathbb{E}N(\nu)\rfloor+1}^{\infty}\mathbb{P}(\dots) \\
&\leq \sum_{n=0}^{\lfloor\epsilon\mathbb{E}N(\nu)\rfloor}\mathbb{P}(N(\nu)=n) + \sum_{n=\lfloor\epsilon\mathbb{E}N(\nu)\rfloor+1}^{\infty}\mathbb{P}(|S(\nu)/N(\nu) - \mathbb{E}X(\nu)| > \delta, N(\nu) = n) \\
&\leq \mathbb{P}(N(\nu)\leq\epsilon\mathbb{E}N(\nu)) + \sum_{n=\lfloor\epsilon\mathbb{E}N(\nu)\rfloor+1}^{\infty}\mathbb{P}(|S_n(\nu)/n - \mathbb{E}X(\nu)| > \delta)\,\mathbb{P}(N(\nu)=n)
\end{aligned}$$

Using Chebyshev's inequality, we find that the second term is bounded from above by

$$\sum_{n=\lfloor\epsilon\mathbb{E}N(\nu)\rfloor+1}^{\infty}\frac{\text{Var}(X(\nu))}{\delta^2 n}\mathbb{P}(N(\nu)=n) \leq \frac{\text{Var}(X(\nu))}{\epsilon\mathbb{E}N(\nu)\delta^2}\sum_{n=\lfloor\epsilon\mathbb{E}N(\nu)\rfloor+1}^{\infty}\mathbb{P}(N(\nu)=n) \leq \frac{\text{Var}(X(\nu))}{\epsilon\mathbb{E}N(\nu)\delta^2}.$$

For every $\xi > 0$ and every $\epsilon > 0$, there exists $\nu_{\epsilon,\xi} > 0$ such that for $\nu > \nu_{\epsilon,\xi}$,

$$\mathbb{P}\left(N(\nu) \leq \epsilon \mathbb{E}N(\nu)\right) \leq \mathbb{P}\left(Z \leq \epsilon\right) + \xi \quad \text{and} \quad \frac{\mathrm{Var}X(\nu)}{\epsilon \mathbb{E}N(\nu)} \leq \xi.$$

The first inequality follows from the fact that the c.d.f. of $N(\nu)/\mathbb{E}N(\nu)$ converges pointwise to that of Z, by hypothesis (i). On the other hand, the second inequality follows immediately from hypothesis (ii). Therefore, for $\nu > \nu_{\epsilon,\xi}$,

$$\mathbb{P}\left(|S(\nu)/N(\nu) - \mathbb{E}X(\nu)| > \delta\right) \leq \mathbb{P}\left(Z \leq \epsilon\right) + \xi + \frac{\xi}{\delta^2}.$$

Take e.g. $\epsilon = \delta$, $\xi = \delta^3$. Then for ν sufficiently large,

$$\mathbb{P}\left(|S(\nu)/N(\nu) - \mathbb{E}X(\nu)| > \delta\right) \leq \mathbb{P}\left(Z \leq \delta\right) + \delta^3 + \delta.$$

Letting $\delta \downarrow 0$ and using the fact that $\mathbb{P}\left(Z = 0\right) = 0$, we obtain (34). Then hypothesis (i) and Slutsky's theorem imply the conclusion. □

D Proof of Lemma 8.2

In this subsection we use the same notation introduced in Section 8, but without writing explicitly the dependence on ν. Denote by π_S the restriction of π to S, i.e. $\pi_S(A) = \pi(A \cap S)$, and define μ_S to be the equilibrium distribution conditioned on S, i.e.

$$\mu_S(A) := \frac{\pi_S(A)}{\pi(S)}. \tag{35}$$

Thanks to one of the characterizations of the total variation distance,

$$\pi(S)\,||\mu_S P^t - \mu_S||_{\mathrm{TV}} = \pi(S) \sum_{\substack{y \in \Omega: \\ \mu_S P^t(y) \geq \mu_S(y)}} \left[\mu_S P^t(y) - \mu_S(y)\right].$$

From (35) it follows that $\pi_S P^t(y) = \pi(S)\mu_S P^t$ and that $\pi_S(y) = \pi(S)\mu_S(y)$. Moreover, $\mu_S P^t(y) \geq \mu_S(y) \iff \pi_S P^t(y) \geq \pi_S(y)$. Therefore

$$\pi(S)\,||\mu_S P^t - \mu_S||_{\mathrm{TV}} = \sum_{\substack{y \in \Omega: \\ \pi_S P^t(y) \geq \pi_S(y)}} \left[\pi_S P^t(y) - \pi_S(y)\right]. \tag{36}$$

Since $\pi_S(x) > 0$ iff $x \in S$ and $\pi_S(x) = \pi(x)$ if $x \in S$,

$$\pi_S P^t(y) = \sum_{x \in \Omega} \pi_S(x) P^t(x,y) = \sum_{x \in S} \pi(x) P^t(x,y) \leq \sum_{x \in \Omega} \pi(x) P^t(x,y) = \pi(y).$$

From this inequality it follows that $\pi_S P^t(y) \leq \pi_S(y)$ for $y \in S$ and trivially $\pi_S P^t(y) \geq \pi_S(y)$ for $y \in S^c$, since the right-hand side is zero. Using these last two inequalities, identity (36) can be rewritten as

$$\pi(S)\,||\mu_S P^t - \mu_S||_{\mathrm{TV}} = \sum_{y \in S^c} \left[\pi_S P^t(y) - \pi_S(y)\right] = \sum_{y \in S^c} \pi_S P^t(y),$$

where the last equation follows from the observation that $\pi_S(y) = 0$ when $y \in S^c$. Furthermore

$$\sum_{y \in S^c} \pi_S P^t(y) = \sum_{y \in S^c} \sum_{x \in S} \pi(x) \mathbb{P}\left(X(t) = y \mid X(0) = x\right) = \sum_{x \in S} \mathbb{P}\left(X(t) \in S^c,\, X(0) = x\right). \tag{37}$$

Define $N_{x \to y}(t)$ as the number of transitions from state x to state y during the time interval $[0, t]$, so that

$$\mathbb{P}\left(X(t) \in S^c,\, X(0) = x\right) \leq \sum_{x' \in S, y' \in S^c} \mathbb{P}\left(N_{x' \to y'}(t) \geq 1,\, X(0) = x\right) \leq \sum_{x' \in S, y' \in S^c} \mathbb{E}\left\{N_{x' \to y'}(t),\, X(0) = x\right\}, \tag{38}$$

and hence, substituting (38) in (37)

$$\begin{aligned}\sum_{y\in S^c}\pi_S P^t(y) &\leq \sum_{x\in S}\sum_{x'\in S, y'\in S^c}\mathbb{E}\left\{N_{x'\to y'}(t),\, X(0)=x\right\}\\ &= \sum_{x'\in S, y'\in S^c}\mathbb{E}\left\{N_{x'\to y'}(t)\right\}\\ &= \sum_{x'\in S, y'\in S^c}\int_{u=0}^{t}\mathbb{P}\left(X(u)=x'\right)\, q(x',y')du\\ &\leq \sum_{x'\in S, y'\in S^c}\pi(x')\, q(x',y')t\\ &= t\, Q(S,S^c).\end{aligned}$$

The second last passage follows from the fact that $\mathbb{P}\left(X(u)=x'\right)=\pi_S P^u(x')\leq \pi_S(x')=\pi(x')$. Then

$$||\mu_S P^t-\mu_S||_{\mathrm{TV}}\leq t\,\Phi(S).$$

Assuming that $\pi(S)\leq r$, $r\in[0,1]$, one has that

$$||\mu_S-\pi||_{\mathrm{TV}}=\max_{A\subset\Omega}|\mu_s(A)-\pi(A)|\geq \pi(S^c)-\mu(S^c)=\pi(S^c)=1-\pi(S)\geq 1-r.$$

The triangular inequality implies that

$$1-r\leq ||\mu_S-\pi||_{\mathrm{TV}}\leq ||\mu_S-\mu_S P^t||_{\mathrm{TV}}+||\mu_S P^t-\pi||_{\mathrm{TV}}.$$

We know that $\bar{d}(t)\leq 2d(t)$ and hence for $t=t_{\mathrm{mix}}(\varepsilon)$, $||\mu_S P^t-\pi||_{\mathrm{TV}}\leq \bar{d}(t_{\mathrm{mix}}(\varepsilon))\leq 2d(t_{\mathrm{mix}}(\varepsilon))=2\varepsilon$. Taking $t=t_{\mathrm{mix}}(\varepsilon)$, it follows that

$$1-r\leq t_{\mathrm{mix}}(\varepsilon)\Phi(S)+2\varepsilon.$$

Rearranging and minimizing over S concludes the proof.